The Open University

M303

Further pure mathematics

Book F

Metric spaces 2

This publication forms part of the Open University module M303 *Further pure mathematics*. Details of this and other Open University modules can be obtained from Student Recruitment, The Open University, PO Box 197, Milton Keynes MK7 6BJ, United Kingdom (tel. +44 (0)300 303 5303; email general-enquiries@open.ac.uk).

Alternatively, you may visit the Open University website at www.open.ac.uk where you can learn more about the wide range of modules and packs offered at all levels by The Open University.

The Open University, Walton Hall, Milton Keynes, MK7 6AA.

First published 2015. Second edition 2016.

Copyright © 2015, 2016 The Open University

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, transmitted or utilised in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without written permission from the publisher or a licence from the Copyright Licensing Agency Ltd. Details of such licences (for reprographic reproduction) may be obtained from the Copyright Licensing Agency Ltd, Saffron House, 6–10 Kirby Street, London EC1N 8TS (website www.cla.co.uk).

Open University materials may also be made available in electronic formats for use by students of the University. All rights, including copyright and related rights and database rights, in electronic materials and their contents are owned by or licensed to The Open University, or otherwise used by The Open University as permitted by applicable law.

In using electronic materials and their contents you agree that your use will be solely for the purposes of following an Open University course of study or otherwise as licensed by The Open University or its assigns.

Except as permitted above you undertake not to copy, store in any medium (including electronic storage or use in a website), distribute, transmit or retransmit, broadcast, modify or show in public such electronic materials in whole or in part without the prior written consent of The Open University or in accordance with the Copyright, Designs and Patents Act 1988.

Edited, designed and typeset by The Open University, using the Open University TEX System.

Printed in the United Kingdom by Halstan & Co. Ltd, Amersham, Bucks.

ISBN 978 1 4730 2041 2

2.1

Contents

Chapter 21 Connectedness **1**

Introduction **3**

1 Homeomorphisms **4**
1.1 Homeomorphic subsets of $\mathbb{R}$: part 1 8

2 Closed and open sets revisited **10**

3 Connectedness **14**
3.1 Using a continuous function to characterise disconnectedness 18
3.2 Connectedness and closure 20
3.3 Components 22
3.4 The components are closed 27
3.5 Totally disconnected spaces 29

4 Connectedness in Euclidean space **30**
4.1 Homeomorphic subsets of $\mathbb{R}$: part 2 35
4.2 Products of connected spaces 36

5 Path-connected spaces **39**
5.1 Defining a path 40
5.2 Examples of path-connected spaces 44
5.3 Products of path-connected spaces 49

6 The topologist's cosine **50**

Learning outcomes **57**

Solutions to exercises **58**

Chapter 22 Compactness **67**

Introduction **69**

1 Two important theorems **70**
1.1 Preliminaries: sequences and subsequences 71

2 Sequential compactness **76**
2.1 Examples of sequentially compact spaces 77
2.2 Subsets, images and products of sequentially compact spaces 79
2.3 Sequentially compact subsets of $\mathbb{R}^n$ 81
2.4 The Generalised Extreme Value Theorem 84

3 Functions and sequential compactness **85**
3.1 Sequentially compact subsets of $\mathbb{R}$ and continuous functions 86
3.2 Sequentially compact subsets of $C[0,1]$ 89
3.3 Proof of Theorem 3.12 95

4 Compact metric spaces **102**
4.1 The Boundedness Theorem revisited 107
4.2 The definition of a compact metric space 108
4.3 Examples of compact spaces 110

5 Sequentially compact and compact are equivalent for metric spaces **111**
5.1 Totally bounded sets 113
5.2 Sequentially compact metric spaces are compact 116
5.3 Compact metric spaces are sequentially compact 118

6 Unions and intersections of compact spaces **120**
6.1 Epilogue 122

Learning outcomes **123**

Solutions to exercises **124**

Chapter 23 Completeness **137**

Introduction **139**

1 Completeness **139**
1.1 Cauchy sequences 141
1.2 Cauchy sequences in $\mathbb{R}$ 145

2 Complete metric spaces **146**
2.1 Criteria for completeness 147
2.2 The completeness of $(\mathbb{R}^n, d^{(n)})$ 149
2.3 Cantor space is complete 150
2.4 $(C[0,1], d_{\max})$ is complete 151
2.5 Completeness is not a homeomorphism invariant 155

3 The Contraction Mapping Theorem **158**
3.1 Fixed points and the Contraction Mapping Theorem 158
3.2 Applications of the Contraction Mapping Theorem 167

4 Completion **171**
4.1 Defining the completion 171
4.2 Constructing the completion of a space 173

Learning outcomes **179**

Solutions to exercises **180**

Chapter 24 Fractals **191**

Introduction **193**

1 Examples of fractals **194**
1.1 Cantor sets 194
1.2 The Sierpiński gasket 199
1.3 The von Koch curve 202

2 Using the Contraction Mapping Theorem **207**
2.1 The Hausdorff metric 207
2.2 Constructing contraction mappings 214
2.3 Iterated function schemes 217
2.4 Proof of Theorem 2.12 222

3 Dimensions of fractals **225**
3.1 Box dimension 225
3.2 The open set condition 230

Conclusion **239**

Learning outcomes **240**

Solutions to exercises **241**

Index **251**

Chapter 21

Connectedness

Introduction

Welcome to the final book of M303. In this book we continue our investigation of metric spaces, but first we pause to take stock. So far you have studied number theory, groups, rings, fields and metric spaces. It is natural to see these as very different topics, but you may be starting to get a sense of an overall similarity. When studying groups, rings and fields, we start by defining a set, together with an operation or two upon that set. We then require that certain rules are satisfied. We often add additional conditions (such as requiring that a group operation is commutative). The integers form a ring, and so we can view number theory as studying one specific ring. A metric space may seem like a very different animal, but the process is the same: we start with a set of elements (such as points on the real line) and add some structure (an idea of 'distance').

$\mathbb{Z}$ is a particularly worthwhile ring to study as there is always a unique ring homomorphism from $\mathbb{Z}$ into any ring with an identity.

In general, algebraists talk of a 'set of elements' and geometers speak of a 'space of points'; but the distinction is somewhat artificial. You may have studied linear algebra and noticed the algebraic concept of addition appearing alongside the geometric notion of a modulus.

There are many mathematical objects which combine algebraic and geometric ideas, such as Lie groups, in which the group operation is required to be, in some sense, continuous.

We can apply varying structures to the same set or space. For example, we can consider $\mathbb{Z}$, the set of integers, as an abelian group, a ring or as a metric space (with an appropriate metric such as the Euclidean metric or the p-adic metric).

When studying mathematical objects, we need to work out what we mean by a different object. For example, we study groups up to isomorphism. We start this chapter by discussing what we mean by 'different' in the context of a metric space.

Our goal in this chapter is to understand what it means to say that a metric space is connected. In doing so, we discuss two different notions of connectedness and discover how they are preserved (or not) under particular transformations of metric spaces known as homeomorphisms.

As a result of these notions of connectedness, we are able to fulfil the first of our promises from Book D, Chapter 13 and prove a version of the Intermediate Value Theorem for continuous functions from a metric space (X, d) to $\mathbb{R}$.

Chapter guide

The key ideas here are those of homeomorphisms, connectedness and path-connectedness. We look at homeomorphisms in Section 1, connectedness in Section 3 and path-connectedness in Section 5. Section 4 looks at connectedness in Euclidean space in particular, and contains a generalisation of the Intermediate Value Theorem. The final section contains an example which illustrates that a connected space need not be path-connected. You should try to understand the ideas underlying this example but you do not need to learn the technical details of the construction.

1 Homeomorphisms

Two groups are isomorphic if we can find a bijection between the two sets which preserves the effect of the group operation. We want a similar idea to describe when two metric spaces (X, d_X) and (Y, d_Y) are 'essentially the same'. Our choice is motivated by a desire to study continuity, in its purest form, and so we use the existence of a bijective (d_X, d_Y)-continuous map, with a (d_Y, d_X)-continuous inverse, between the two spaces to define our idea of equivalence.

Definition 1.1 *Homeomorphism*

Let (X, d) and (Y, e) be two metric spaces. Then (X, d) and (Y, e) are **homeomorphic** metric spaces if there exists a bijection $f\colon X \to Y$ such that $f\colon X \to Y$ is (d, e)-continuous and $f^{-1}\colon Y \to X$ is (e, d)-continuous. We say that f is a **homeomorphism**.

Topology (a generalisation of metric spaces mentioned in Book D, Chapter 16) is often called 'rubber sheet geometry' for this reason.

Intuitively, two spaces are homeomorphic if one can be continually deformed into the other. We can twist and stretch our spaces, but we cannot create holes (which would create a discontinuity).

The Composition and Restriction Rules for continuous functions generalise as expected to homeomorphisms. In particular, we have the following result.

Proposition 1.2 *Restriction Rule for homeomorphisms*

Let $h\colon X \to Y$ be a homeomorphism between two metric spaces (X, d) and (Y, e), and let $X' \subseteq X$. Then $h|_{X'}$, the restriction of h to X', gives a homeomorphism from X' to the image $h(X') \subseteq Y$ so that $(X', d_{X'})$ and $(h(X'), e_{h(X')})$ are homeomorphic metric spaces.

The induced metric, d_A, for a subset A of a metric space (X, d) is given by $d_A(a, b) = d(a, b)$ for $a, b \in A$.

Proof We need to show that $h|_{X'}$ is $(d_{X'}, e_{h(X')})$-continuous, is a bijection, and has a $(e_{h(X')}, d_{X'})$-continuous inverse. As h is (d, e)-continuous and h^{-1} is (e, d)-continuous, the Restriction Rule for continuous functions (Book D, Chapter 15, Theorem 1.6) shows that $h|_{X'}$ is $(d_{X'}, e_{h(X')})$-continuous and $h^{-1}|_{h(X')}$ is $(e_{h(X')}, d_{X'})$-continuous. Clearly $h|_{X'}$ and $h^{-1}|_{h(X')}$ are one–one (since h and h^{-1} are) and, by definition, $h|_{X'}$ is onto its image and $h^{-1}|_{h(X')}$ is onto its preimage. ■

Example 1.3

Consider the sets $X = (0, 1)$ and $Y = (0, 5)$ (each with the metric induced by the Euclidean metric on $\mathbb{R}$). Then the map $f\colon x \mapsto 5x$ gives a homeomorphism from X to Y, and so $(X, d_X^{(1)})$ and $(Y, d_Y^{(1)})$ are homeomorphic metric spaces.

Example 1.4

Consider the sets $X = [1, 2]$ and $Y = [-10, 10]$ (each with the metric induced by the Euclidean metric on $\mathbb{R}$). Then the map $g\colon X \to Y$ given by $g\colon x \mapsto -10 + 20(x - 1)$ is a homeomorphism from X to Y, and so $(X, d_X^{(1)})$ and $(Y, d_Y^{(1)})$ are homeomorphic metric spaces.

Example 1.5

Let S be the square $\{(x, y) \in \mathbb{R}^2 \colon 0 \leq x \leq 1 \text{ and } 0 \leq y \leq 1\}$ and R be the rectangle $\{(x, y) \in \mathbb{R}^2 \colon 0 \leq x \leq 2 \text{ and } 0 \leq y \leq 1\}$, each with the metric induced by the Euclidean metric on $\mathbb{R}^2$. The map $h\colon S \to R$ given by $h\colon (x, y) \mapsto (2x, y)$ gives a homeomorphism from S to R, and so $(S, d_S^{(2)})$ and $(R, d_R^{(2)})$ are homeomorphic metric spaces.

Exercise 1.1

Let $A = [0, 1]$ and $B = [0, n]$ for some $n \in \mathbb{N}$, each with the metric induced by the Euclidean metric on $\mathbb{R}$. Show that $(A, d_A^{(1)})$ and $(B, d_B^{(1)})$ are homeomorphic metric spaces by writing down a specific homeomorphism from A to B. Can n be any real number?

Exercise 1.2

Let $X = [a, b]$ and $Y = [c, d]$, each with the metric induced by the Euclidean metric on $\mathbb{R}$. By writing down an explicit homeomorphism between X and Y, show that any two closed intervals $[a, b]$ and $[c, d]$ with $a \neq b$ and $c \neq d$, and each with the Euclidean metric, are homeomorphic. Does it matter if $a < b$ and $c > d$?

Exercise 1.3

Let $X = [a, b)$ and $Y = [c, d)$, each with the metric induced by the Euclidean metric on $\mathbb{R}$. Show that if $a \neq b$ and $c \neq d$ then X and Y are homeomorphic metric spaces. Is the same true for $(a, b]$ and $[c, d)$?

Instead of writing down homeomorphisms explicitly (especially when working in Euclidean space), we often refer to the 'obvious' map. However, it is important to be sure that 'the obvious map' is a continuous bijection with a continuous inverse!

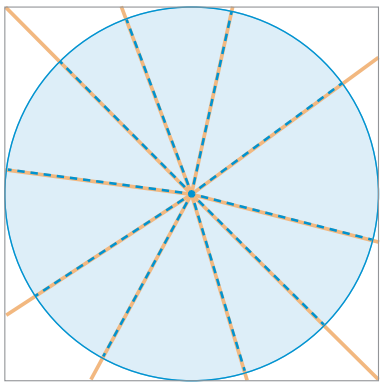

Figure 1.1

Example 1.6

The closed disc, $D = \{(x, y) \in \mathbb{R}^2 : x^2 + y^2 \leq 1\}$, and the square $S = \{(x, y) \in \mathbb{R}^2 : -1 \leq x \leq 1 \text{ and } -1 \leq y \leq 1\}$, each with the metric induced by the Euclidean metric on $\mathbb{R}^2$, are homeomorphic metric spaces. The homeomorphism is given by mapping the radial lines of the disc onto the corresponding radial lines of the square; the function 'stretches' the radial lines of the disc so that they map in a one–one way onto the radial lines of the square (see Figure 1.1).

Example 1.7

The unit ball in $\mathbb{R}^3$, $B_{d(3)}[(0, 0, 0), 1]$ and the cube $\{(x, y, z) \in \mathbb{R}^3 : -1 \leq x \leq 1 \text{ and } -1 \leq y \leq 1 \text{ and } -1 \leq z \leq 1\}$, each with the metric induced by the Euclidean metric on $\mathbb{R}^3$, are homeomorphic metric spaces. The homeomorphism is given by a three-dimensional version of the mapping shown in Figure 1.1.

Exercise 1.4

Convince yourself that a coffee cup is homeomorphic to a doughnut (when each are viewed as subsets of $\mathbb{R}^3$ and given the metric induced by the Euclidean metric on $\mathbb{R}^3$).

Once we start to think about homeomorphisms, we sometimes find that the closed-set characterisation of continuity, given in Theorem 7.1 of Chapter 16, can be more useful than the sequential definition. We restate this characterisation here for reference.

Theorem 1.8

Let (X, d) and (Y, e) be metric spaces and suppose that $f : X \to Y$. Then f is (d, e)-**continuous** on X if, and only if, for each e-closed set $E \subseteq Y$ $f^{-1}(E)$ is d-closed.

Remark

We could, of course, phrase this in terms of open sets rather than closed sets (using Book D, Chapter 16, Corollary 7.2).

We use this characterisation of continuity in the following exercise to demonstrate an unusual property of discrete metric spaces.

Worked Exercise 1.9

Let (X, d_0) be any space with the discrete metric. Let (Y, d) be any metric space and $f\colon X \to Y$ be any function. Show that f is (d_0, d)-continuous.

Solution

From Worked Exercise 1.11 of Chapter 16 we know that all subsets of a discrete metric space are d_0-closed. Therefore for any d-closed set $A \subseteq Y$, we must have that $f^{-1}(A)$ is d_0-closed in X, so f is (d_0, d)-continuous by Theorem 7.1 of Chapter 16.

We have now proved the following.

Proposition 1.10

Let (X, d_0) be any space with the discrete metric. Let (Y, d) be any metric space and $f\colon X \to Y$ be any function. Then f is (d_0, d)-continuous.

Discrete metric spaces have rather unintuitive properties (and so are a good source of counter-examples to conjectures that might seem 'obviously true').

We defined equivalent metrics for the same underlying set in Definition 3.5 of Chapter 15. It is natural to wonder whether two spaces with equivalent metrics are homeomorphic. We ask you to determine this in the next worked exercise.

Worked Exercise 1.11

Consider a set X with two equivalent metrics d and d'. Show that (X, d) and (X, d') are homeomorphic metric spaces.

Solution

We need to find a (d, d')-continuous bijective map, $f\colon X \to X$ with a (d', d)-continuous inverse. Consider the identity map $\mathrm{Id}\colon X \to X$. This is clearly bijective and (d, d)-continuous. Therefore by Theorem 3.8 of Chapter 15 it is (d, d')-continuous. The same result shows that the inverse, $\mathrm{Id}^{-1}\colon X \to X$, is (d', d)-continuous. Therefore (X, d) and (X, d') are homeomorphic metric spaces.

$\mathrm{Id}(x) = x$ for each $x \in X$.

Worked Exercise 1.11 shows that we cannot tell the difference between metric spaces with equivalent metrics when we are working up to homeomorphism. It turns out that many properties of metric spaces are preserved by homeomorphism. By this we mean that if X and Y are homeomorphic spaces, and X has one of these properties, then Y must also have the property. Properties preserved by homeomorphism are called *topological invariants*.

The collection of all the open sets in a metric space is called the topology on the space (Book D, Chapter 16, Definition 7.3).

Definition 1.12 *Topological invariant*

We say that a property of metric spaces is a **topological invariant** if it is preserved by homeomorphisms. That is, if two spaces are homeomorphic and one has the property, then so must the other.

Connectedness, intuitively, tells us whether or not a space is 'in one piece'. Compactness is a concept which gives us an idea of 'size' for a metric space. Completeness, in some sense, tells us whether a space has any 'missing points'.

We will show that the two properties explored in the first two chapters of this book, namely connectedness and compactness, are both topological invariants but that the last of the 'three Cs', completeness, is not.

1.1 Homeomorphic subsets of $\mathbb{R}$: part 1

We finish this section by taking a first look at which intervals of $\mathbb{R}$ are homeomorphic.

First we note that if a set $A \subseteq \mathbb{R}$ is an interval, then whenever $a, b \in A$ and $c \in \mathbb{R}$ with $a < c < b$, we have $c \in A$. In particular, we can view $\mathbb{R}$ as an interval by thinking of it as $(-\infty, \infty)$. The empty set, $\emptyset = (a, a)$ and a single point $\{a\} = [a, a]$ are also intervals.

In this section we assume that all intervals have the metric induced by the Euclidean metric on $\mathbb{R}$, which we write as d. In Exercise 1.2 we used the map $f_I : x \mapsto c + \frac{d-c}{b-a}(x-a)$ to show that any two closed intervals of $\mathbb{R}$ are homeomorphic. The restriction of this map to the open interval (a, b) shows that any two bounded open intervals are homeomorphic. We also used this map in Exercise 1.3 to show that any two bounded half-open intervals are homeomorphic. In other words, any two bounded intervals of the same 'type' are homeomorphic.

A *bounded* interval is one of the form $[a, b], [a, b), (a, b]$ or (a, b) for $a, b \in \mathbb{R}$.

In the next exercise we ask you to show that the map f_I does not give a homeomorphism between bounded intervals of different types.

Exercise 1.5

Show that $f : x \mapsto c + \frac{d-c}{b-a}(x-a)$ is not a homeomorphism from $(a, b]$ to (c, d).

Of course the fact that f_I does not give us an explicit homeomorphism between intervals of different types does not prove that no such map exists. Showing that two spaces are not homeomorphic can be difficult, and mathematicians often use properties that are preserved by homeomorphism to distinguish between such spaces. In Subsection 4.1 we use a topological invariant, *connectedness*, to prove that intervals of different types are not homeomorphic.

Sometimes we can show that two sets are not homeomorphic without resorting to topological invariants. We ask you to do this in the next exercise.

Exercise 1.6

Show that the metric space $(\mathbb{R}, d)$, where d is the Euclidean metric, and the set $\{1, 2\}$ with the discrete metric, are not homeomorphic metric spaces.

It may seem intuitively obvious that any two intervals of the same type are homeomorphic. What is more surprising is that the whole of $\mathbb{R}$ itself (with the Euclidean metric) and the interval $(0, 1)$ (with the metric induced by the Euclidean metric on $\mathbb{R}$) are homeomorphic metric spaces.

Worked Exercise 1.13

Show that the metric space $(\mathbb{R}, d^{(1)})$ and the metric space $\big((0,1), d^{(1)}_{(0,1)}\big)$ are homeomorphic metric spaces.

Solution

To see this we consider the function $\phi\colon \mathbb{R} \to (0,1)$ given by $\phi(x) = \frac{1}{2} + \frac{1}{\pi}\arctan(x)$ (taking the principal value of arctan). This is the sum of basic continuous functions (so is $(d^{(1)}, d^{(1)}_{(0,1)})$-continuous) and has a $(d^{(1)}_{(0,1)}, d^{(1)})$-continuous inverse, namely $\tan\big(\pi(x - 1/2)\big)$. The fact that the inverse is defined on $(0,1)$ shows that ϕ is onto, and standard properties of arctan show that it is one–one.

Remark

Don't worry if the map ϕ appeared in the above discussion like a rabbit out of a hat. Maps like this do not usually immediately spring to mind. They are constructed, often after considerable effort, by thinking about the kind of map that might work and then modifying it. In this case, we want a map that squashes all of the real line into $(0,1)$. We know that arctan squashes things up, so we use that and then fiddle around to get the scaling right.

Exercise 1.7

Show that $(\mathbb{R}, d^{(1)})$ and $\left((a,\infty), d^{(1)}_{(a,\infty)}\right)$, where a is any real number, are homeomorphic metric spaces.

To summarise, we know that any two bounded intervals of the same 'type' are homeomorphic, and that $\mathbb{R}$ is homeomorphic to the $d^{(1)}$-open interval $(0,1)$ (and so, since the composition of two homeomorphisms is also a homeomorphism, we know that $\mathbb{R}$ is homeomorphic to any non-empty $d^{(1)}$-open interval). The map $f\colon x \mapsto \tan(x)$ restricted to the interval $[0, \frac{\pi}{2})$ can be used to give a homeomorphism between $[0, \frac{\pi}{2})$ and $[0,\infty)$. That, combined with Exercise 1.3, and the fact that the composition of two homeomorphisms is a homeomorphism, shows that any two half-open intervals are homeomorphic.

Therefore we have now identified five classes of intervals:

- the empty set $\emptyset$
- the singleton sets $\{a\} = [a,a]$
- the open intervals (a,b) for $a < b$, (a,∞) and $(-\infty,a)$ for $a \in \mathbb{R}$, and $\mathbb{R}$
- the closed intervals $[a,b]$ for $a < b$
- the intervals of the form $[a,b)$ and $(a,b]$ for $a < b$ and $[a,\infty)$ and $(-\infty,a]$.

Intervals in the same class are homeomorphic. We do not yet know whether intervals in different classes are not homeomorphic. We shall investigate this question in Subsection 4.1.

2 Closed and open sets revisited

Recall from Chapter 16 that a set, E, is d-closed for a metric d if every d-convergent sequence (x_n), with terms $x_n \in E$, has its limit, x, in E; and that a d-open set is the complement of a d-closed set.

When examining a *subspace* E of a metric space (X,d) we must be a little careful. Consider the interval $(0,1)$. Stop for a moment (not for too long!) and think about whether this interval is closed ....

You may have decided that this interval is not closed – we tend to refer to intervals like this as 'open intervals'. The interval $(0,1)$ is certainly not $d^{(1)}$-closed as a subset of $\mathbb{R}$ with the Euclidean metric $d^{(1)}$, but is it $d^{(1)}_{(0,1)}$-closed as a subset of $(0,1)$ with the induced metric $d^{(1)}_{(0,1)}$? In other words, what happens if we consider the set $(0,1)$, with the subspace metric, $d^{(1)}_{(0,1)}$, as our whole space? We know from Chapter 16, Lemma 1.2 that the whole space, X, is always a d-closed subset of a metric space (X,d), so if we consider the set $(0,1)$ as a subset of the metric space

$\left((0,1), d^{(1)}_{(0,1)}\right)$ then it is $d^{(1)}_{(0,1)}$-closed. So, whether or not a particular set is closed depends on 'where this set is' (and of course on which metric we are using). The interval $(0,1)$ is not $d^{(1)}$-closed as a subset of $\left(\mathbb{R}, d^{(1)}\right)$, but is $d^{(1)}_{(0,1)}$-closed as a subset of $\left((0,1), d^{(1)}_{(0,1)}\right)$.

It is natural to ask whether there is a space, (X,d), containing $(0,1)$ as a *proper* subset in which $(0,1)$ is d-closed. In order to answer this, it is helpful to use the characterisation of closed sets introduced in Theorem 3.12 of Chapter 16, a characterisation which turns out to be so useful in what follows that we repeat it here for reference.

Let (X,d) be a metric space and let $A \subseteq X$. A set $E \subseteq A$ is d_A-closed if, and only if, there is a set F in X that is d-closed and for which

$$E = A \cap F.$$

In other words, if $A \subseteq X$ is a subset in a metric space (X,d) then the d_A-closed sets of A are the intersections of A with the d-closed sets of X.

Exercise 2.1

Find a space, (X,d), containing $(0,1)$ as a proper subset in which $(0,1)$ is d-closed.

As you may expect there is an equivalent characterisation for the open sets of a subspace.

Theorem 2.1

Let (X,d) be a metric space and let $A \subseteq X$.

A set $U \subseteq A$ is d_A-open if, and only if, there is a set E in X that is d-open and for which

$$U = A \cap E.$$

Proof

The set U is d_A-open if, and only if, the complement of U in A, which we write as V, is d_A-closed. Therefore U is d_A-open if, and only if, $V = \{a \in A : a \notin U\}$ is d_A-closed. By Theorem 3.12 of Chapter 16, we know that V is d_A-closed if, and only if, there is a d-closed set F in X with $V = A \cap F$; see Figure 2.1. Since F is d-closed, $X - F$ is d-open.

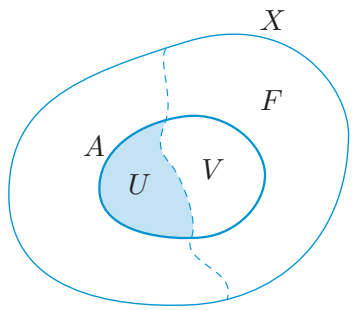

Figure 2.1

Now

$$\begin{aligned} A \cap (X - F) &= \{a \in A \colon a \notin F\} \\ &= \{a \in X \colon a \in A \text{ and } a \notin F\} \\ &= \{a \in X \colon a \in A \text{ and } a \in (A - F)\} \\ &= \{a \in X \colon a \in A \text{ and } a \notin V\} \\ &= \{a \in A \colon a \notin V\} \\ &= U \end{aligned}$$

and so U is the intersection of A with the d-open set $X - F$. ■

When working with subsets of subspaces, it can be necessary to carefully 'unpack' the situation so as to be very clear about exactly what is happening. (This is, in fact, true of many areas of pure mathematics.) The following worked example demonstrates this 'unpacking'.

Worked Exercise 2.2

Let $A = (-\infty, 0) \cup (1, \infty)$ with the metric induced by $d^{(1)}$, the Euclidean metric defined on $\mathbb{R}$. Show that the subset $E \subseteq A$, given by $E = \{x \colon -7 \leq x < 0\}$ is d_A-closed.

Solution

Unpacking the definition, we are looking at $E = [-7, 0) \subseteq A \subseteq \mathbb{R}$. The space $A = (-\infty, 0) \cup (1, \infty)$ has the induced metric $d_A^{(1)}$. Therefore the $d_A^{(1)}$-closed sets in A are the intersections of A with the $d^{(1)}$-closed sets in $\mathbb{R}$. This means that the subset $E = [-7, 0) \subseteq A$ is $d_A^{(1)}$-closed as a set in A if there is a $d^{(1)}$-closed subset F of $\mathbb{R}$ such that $[-7, 0) = F \cap A$. Any interval of the form $[a, b]$ is a $d^{(1)}$-closed subset of $\mathbb{R}$, so if we find an interval $[a, b]$ with $[-7, 0) = [a, b] \cap \big((-\infty, 0) \cup (1, \infty)\big)$ then we are done. The interval $[-7, 0]$ is clearly such an interval. (In fact, any interval of the form $[-7, b]$ with $0 \leq b \leq 1$ will suffice.)

Of course you do not always have to write down the details as fully as we have in the example above. However, stopping and thinking about which space you are 'in', and exactly which metric you are using (especially if you are looking at subsets of a space with a subspace metric) is often helpful.

The next exercises give you practice in recognising $d_X^{(1)}$-closed and $d_X^{(1)}$-open subsets of X, where $X \subseteq \mathbb{R}$ and $(X, d_X^{(1)})$ is a metric space with the metric induced from the Euclidean metric on $\mathbb{R}$.

Exercise 2.2

(a) Show that the subset $\{x \in \mathbb{R}: -\infty < x < 0\} = (-\infty, 0)$ is simultaneously d-closed and d-open when considered as a subset of (X, d) where $X = (-\infty, 0) \cup (1, \infty)$ and d is the metric induced by $d^{(1)}$, the Euclidean metric on $\mathbb{R}$.

(b) Show that the subset $\{x \in \mathbb{R}: -\infty < x \leq 0\} = (-\infty, 0]$ is simultaneously d-closed and d-open when considered as a subset of (X, d) where $X = (-\infty, 0] \cup (1, \infty)$ and d is the metric induced by $d^{(1)}$, the Euclidean metric on $\mathbb{R}$.

(c) Show that the subset $\{0\}$ is simultaneously d-closed and d-open when considered as a subset of (X, d) where $X = \{0\} \cup (3, 4]$ and d is the metric induced by $d^{(1)}$, the Euclidean metric on $\mathbb{R}$.

(d) Show that the subset $(0, 1)$ is simultaneously d-closed and d-open when considered as a subset of $\big((\mathbb{R} - \mathbb{Z}), d\big)$ where d is the metric induced by $d^{(1)}$, the Euclidean metric on $\mathbb{R}$.

Summary

We have two equivalent characterisations of continuity, which we note here for reference and comparison. We also take the opportunity to summarise some facts concerning the word 'closed'.

Let (X, d) and (Y, e) be metric spaces and let $f: X \to Y$ be a function.

- *The sequential characterisation of continuity*
  The function f is (d, e)-**continuous** at $a \in X$ if, and only if, whenever (a_k) is a sequence in X for which $a_k \xrightarrow{d} a$ as $k \to \infty$, then the sequence $f(a_k) \xrightarrow{e} f(a)$ as $k \to \infty$.
- *The 'preimage of closed is closed' characterisation*
  The function f is (d, e)-**continuous** if, and only if, $f^{-1}(U)$ is a d-closed set in (X, d) whenever U is an e-closed set in (Y, e).

Let (X, d) be a metric space and let A be a subset of X.

- The set A is a d-closed subset of X if, and only if, every d-convergent sequence with all its terms in A has its limit in A.
- The set A is d-closed if, and only if, the complement, $A^c = X - A$, is d-open (that is, if A^c contains a d-open ball around each of its points).

Let (X, d) be a metric space. Let E be a subset of A where A is a subset of X, and let d_A be the metric on A induced by the metric d on X.

- The subset E of A is d_A-closed considered as a subset of (A, d_A) if, and only if, there is a d-closed subset F of X with $E = A \cap F$.

3 Connectedness

We are now in a position to think about what we mean by 'connected'. We have an intuitive idea of what we mean for subsets in Euclidean space. We want a space such as $(0,1) \cup (2,3)$, which naturally falls into two pieces, to be disconnected. However, we want to capture the idea of 'falling apart' and so exclude cases such as $(0,1] \cup (1,2) = (0,2)$ where we have artificially 'cut' the space into two pieces. Let us examine the difference between these two cases.

Pinning down exactly what we want an intuitive idea (like 'fall apart') to mean, and translating it into the precise language of mathematics, is a very important (and often difficult) part of mathematics.

First consider $(0,1) \cup (2,3)$. This has an obvious 'gap' in the middle. Moving the second interval closer to $(0,1)$, giving $(0,1) \cup (\frac{3}{2}, \frac{5}{2})$, retains this 'gap'. Imagine moving the second interval closer and closer to the first. How small should we allow this gap to become before deciding that the space no longer 'falls apart'? Write down the space with the 'smallest' gap ....

... You may have written $(0,1) \cup (1,2)$. We want our idea of connected to encapsulate the way that $A = (0,1) \cup (1,2)$ 'falls apart' whereas $B = (0,1] \cup (1,2) = (0,2)$ does not. You may have spotted a key difference between these two cases. The set $(0,1) \cup (1,2)$ falls into two subsets, each of which is simultaneously $d_A^{(1)}$-closed and $d_A^{(1)}$-open, whereas neither $(0,1]$ nor $(1,2)$ is simultaneously $d_B^{(1)}$-closed and $d_B^{(1)}$-open. We know that for any space metric space (X, d), the whole space, X itself, is always simultaneously d-closed and d-open. It therefore seems that being simultaneously closed and open in this way is a rather special quality of a space which, in some sense, makes it into a 'whole piece' in its own right (rather than being merely a 'bit' of another space). We use this idea in our definition of connected.

Definition 3.1 *Disconnection*

Let (X, d) be a metric space. A **disconnection** $\{U, V\}$ of X is a pair of disjoint non-empty subsets, U and V, each of which is simultaneously d-closed and d-open, with $X = U \cup V$. We say that the space (X, d) is **disconnected** if X has a disconnection. It is **connected** if X has no disconnection.

When we wish to emphasise the metric, we say *d-disconnection*, *d-connected* and *d-disconnected*.

Since U and V are required to be non-empty, neither U nor V can be X. Thus U and V must be proper non-empty subsets of X.

Recall that $A \subseteq X$ is a *proper* subset of X if $A \neq X$.

Exercise 3.1

Write down a disconnection of each of the following.

(a) The metric space (X, d_X) where $X = (-\infty, 0) \cup (1, \infty)$ and d_X is the metric induced by the Euclidean metric on $\mathbb{R}$.

(b) The metric space (X, d_X) where $X = (-\infty, 0] \cup (1, \infty)$ and d_X is the metric induced by the Euclidean metric on $\mathbb{R}$.

(c) The metric space (X, d_X) where $X = \{0\} \cup (3, 4]$ and d_X is the metric induced by the Euclidean metric on $\mathbb{R}$.

We now ask you to show that some sets are connected.

Exercise 3.2

Show that the empty set is a connected space regardless of the choice of metric.

Exercise 3.3

Let X consist of a single element $X = \{x\}$. Show that (X, d) is connected (where d is the unique metric defined on the space and given by $d(x, x) = 0$).

The definition of a disconnection involves two subsets. However, the existence of just one non-empty, closed and open, proper subset is enough to guarantee that a space is disconnected.

Theorem 3.2

The metric space (X, d) is connected if, and only if, the only subsets of X which are simultaneously d-closed and d-open are $\emptyset$ and X.

Proof By definition, we know that if X has a disconnection $U \cup V$ then U and V are non-empty, proper subsets which are each simultaneously d-closed and d-open. Hence if the only d-closed and d-open subsets of X are $\emptyset$ and X, then we cannot find appropriate sets U and V and so (X, d) is connected.

Suppose now that U is a non-empty and proper d-closed and d-open subset of X. We must show that (X, d) is disconnected. Define $V = U^c$. Since U is d-closed, V is d-open and since U is d-open, V is d-closed. Since U is proper, V is non-empty and since V is proper, U is non-empty. Hence U

and V are disjoint non-empty d-closed and d-open sets and their union is X. Thus $\{U, V\}$ is a disconnection of X, and so (X, d) is disconnected. ■

Worked Exercise 3.3 *Cantor space*

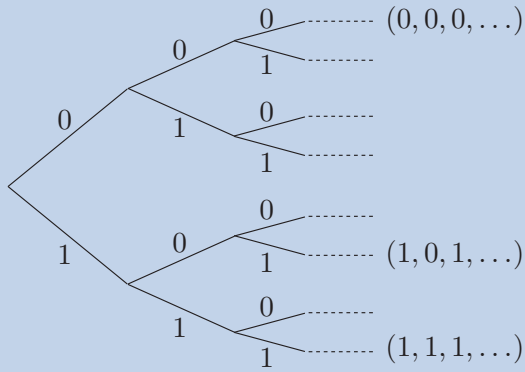

Figure 3.1

Cantor space, $\mathbf{C}$, and Cantor distance, $d_{\mathbf{C}}$, are defined in Book D, Chapter 15, Section 2.

Show that the Cantor space $\mathbf{C}$ is $d_{\mathbf{C}}$-disconnected.

Solution

Recall from Chapter 15 that the Cantor space $\mathbf{C}$ consists of all infinite sequences of zeros and ones. In Chapter 16 we showed that the open balls in $\mathbf{C}$, which are $d_{\mathbf{C}}$-open, are also $d_{\mathbf{C}}$-closed sets. Therefore by Theorem 3.2 Cantor space is disconnected.

Exercise 3.4 *The discrete metric*

Let X be a set with at least two elements, and let d_0 be the discrete metric on X. Show that (X, d_0) is disconnected.

Disconnectedness can also be described in terms of closed (or open) sets.

Exercise 3.5

Prove that a metric space (X, d) is disconnected if, and only if, X can be expressed as the disjoint union of two non-empty d-closed sets.

A similar proof shows that a metric space (X, d) is disconnected if, and only if, X can be expressed as the disjoint union of two non-empty d-open sets. We have the following result.

Proposition 3.4

The following statements are equivalent for a metric space (X, d).

(a) (X, d) is d-disconnected.

(b) X can be expressed as the disjoint union of two non-empty d-closed sets.

(c) X can be expressed as the disjoint union of two non-empty d-open sets.

Some texts define a metric space (X, d) as being disconnected if it can be expressed as the disjoint union of two non-empty d-open sets.

Worked Exercise 3.5

Determine whether $(\mathbb{Q}, d_{\mathbb{Q}})$ is connected, when $d_{\mathbb{Q}}$ is the induced metric on $\mathbb{Q}$ inherited from $\mathbb{R}$ with the Euclidean metric.

Solution

We look for a way to divide $\mathbb{Q}$ into two disjoint, $d_{\mathbb{Q}}$-closed subsets of $\mathbb{Q}$. The set $\mathbb{Q}$ has the metric, $d_{\mathbb{Q}}$, induced by the Euclidean metric on $\mathbb{R}$. Therefore, by Book D, Chapter 16, Theorem 3.12, we know a subset of $\mathbb{Q}$ is $d_{\mathbb{Q}}$-closed if, and only if, it is the intersection of a $d^{(1)}$-closed subset of $\mathbb{R}$ with $\mathbb{Q}$.

Let x be an irrational number, such as $\sqrt{2}$. Then $E = (-\infty, \sqrt{2}] \cap \mathbb{Q}$ is the intersection of a $d^{(1)}$-closed set in $\mathbb{R}$ with $\mathbb{Q}$, and so is $d_{\mathbb{Q}}$-closed. Similarly, $F = [\sqrt{2}, \infty) \cap \mathbb{Q}$ is $d_{\mathbb{Q}}$-closed. These sets are disjoint (since $\sqrt{2} \notin \mathbb{Q}$) and $\mathbb{Q} = E \cup F$. Therefore, by Proposition 3.4, $(\mathbb{Q}, d_{\mathbb{Q}})$ is disconnected.

We ask you to try a similar exercise.

Exercise 3.6

Determine whether $(\mathbb{R} - \mathbb{Q}, d_{\mathbb{R}-\mathbb{Q}})$ is connected, when $d_{\mathbb{R}-\mathbb{Q}}$ is the metric on $\mathbb{R} - \mathbb{Q}$ induced by the Euclidean metric on $\mathbb{R}$.

As we would expect, homeomorphic spaces come in the same number of 'pieces', and connectedness is a topological property.

Lemma 3.6

Let (X, d) and (Y, d') be metric spaces, with $f\colon X \to Y$ a (d, d')-continuous map. If (X, d) is d-connected then $\left(f(X), d'_{f(X)}\right)$ is a $d'_{f(X)}$-connected metric space. In other words, the continuous image of a connected metric space is connected.

Proof Let (X, d) be a connected metric space and let $f\colon X \to Y$ be a (d, d')-continuous map from X to Y. Assume for a contradiction that $E \cup F = f(X)$ is a $d'_{f(X)}$-disconnection of the image $f(X)$. Then E and F are each proper $d'_{f(X)}$-closed non-empty subsets of $f(X)$. Consider the preimages $E_X = f^{-1}(E)$ and $F_X = f^{-1}(F)$. Clearly $X = E_X \cup F_X$ and each of the sets E_X and F_X is non-empty. Also $E_X \cap F_X = f^{-1}(E) \cap f^{-1}(F) = f^{-1}(E \cap F)$. Therefore since E and F are disjoint we must have E_X and F_X disjoint. The function f is (d, d')-continuous, so the preimage of a $d'_{f(X)}$-closed subset of Y is a d-closed subset of X (by Book D, Chapter 16, Theorem 7.1). Therefore both E_X and F_X are d-closed and so form a d-disconnection of X. This contradicts the fact that X is d-connected, so the assumption is false and $f(X)$ is d-connected. ■

Corollary 3.7

Connectedness is preserved by homeomorphism.

Proof Let (X, d) and (Y, e) be homeomorphic metric spaces and let $f\colon X \to Y$ be a homeomorphism between them. The map $f\colon X \to Y$ is onto and a (d, e)-continuous map. Therefore if X is d-connected then Lemma 3.6 implies that Y must be e-connected. Also $f^{-1}\colon Y \to X$ is an onto (e, d)-continuous map from Y to X so, by the same lemma, X is d-connected if Y is e-connected. Therefore X is d-connected if, and only if, Y is e-connected. ■

The converse of Lemma 3.6 is not true. We ask you to show this in the next exercise.

Exercise 3.7

Give an example of two metric spaces (X, d) and (Y, e) and an onto (d, e)-continuous function $f\colon X \to Y$ such that (X, d) is disconnected but (Y, e) is connected.

3.1 Using a continuous function to characterise disconnectedness

We now know that the continuous image of a connected space is connected. This means that if the continuous image of a space is disconnected, then that space must be disconnected. In particular, we can characterise connectedness by considering continuous maps from a space to a disconnected space. Arguably the 'simplest' disconnected space is $\big(\{0, 1\}, d_0\big)$ where d_0 is the discrete metric. We use this space in our characterisation.

There is nothing special about the numbers 0 and 1: we could use any set with two elements here.

Definition 3.8 *Characteristic function*

Let (X, d) be a metric space. If $U \subseteq X$, then the **characteristic function** of U is the function $\chi_U \colon X \to \{0, 1\}$ (Figure 3.2) defined by

$$\chi_U(x) = \begin{cases} 1, & \text{if } x \in U, \\ 0, & \text{if } x \in U^c. \end{cases}$$

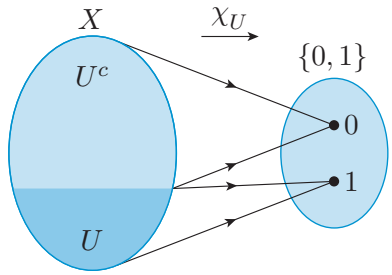

Figure 3.2

In other words, $\chi_U(x)$ is simply the function which takes all elements of U to 1 and everything else to 0. In the following exercises we ask you to investigate what it means for this function to be continuous.

Exercise 3.8

Let (X, d) be a metric space and let $U \subseteq X$. Show that if the characteristic function $\chi_U \colon X \to \{0, 1\}$ is (d, d_0)-continuous, where d_0 is the discrete metric, then U is both d-closed and d-open.

Exercise 3.9

Let (X, d) be a metric space and let $U \subseteq X$. Show that if U is both d-closed and d-open then the characteristic function $\chi_U \colon X \to \{0, 1\}$ is (d, d_0)-continuous (where d_0 is the discrete metric on $\{0, 1\}$).

More generally, let (X, d) be a metric space. Suppose that we can find a (d, d_0)-continuous function $f \colon X \to Y$ where $Y = \{0, 1\}$ and d_0 is the discrete metric. Then

$$U = f^{-1}(\{1\}) \quad \text{and} \quad U^c = f^{-1}(\{0\})$$

are disjoint closed and open sets. Thus X is disconnected (unless one of U or U^c is the empty set, in which case f is a constant function).

Here, f could be one of two constant functions: $f(x) = 0$ for all $x \in X$, or $f(x) = 1$ for all $x \in X$.

This gives the following useful characterisation of connectedness.

Theorem 3.9

Let (X, d) be a metric space. Then X is d-connected if, and only if, every (d, d_0)-continuous function $f \colon X \to (Y, d_0)$, where $Y = \{0, 1\}$ and d_0 is the discrete metric, is constant.

Remark

One consequence of this theorem is that if X is connected, then no function from X to $\{0, 1\}$ can be both continuous and *onto*.

3.2 Connectedness and closure

Many properties are not preserved by closure: the closure of the rational numbers is the real numbers and so countability is not preserved by closure

In this section we use the characterisation of connectedness given by Theorem 3.9 to show that connectedness is preserved by the closure operation. In other words, if (X, d) is a metric space and $A \subseteq X$ is a connected subset of X, then $\mathrm{Cl}_{(X,d)}(A)$ is a connected subset of X. Intuitively, when forming the closure of a set we are simply adding in all the limit points of any sequences in the set. Therefore we would expect connectedness to be preserved since all the points that we add in are 'very close' to the set.

In fact, we prove a stronger result and show that if A is connected then every subset of $\mathrm{Cl}_{(X,d)}(A)$ containing A is connected. In order to do this, we first need a lemma which says that if (X, d) is a metric space and A is a subset of X then any non-empty d-open subset, V, of $\mathrm{Cl}_{(X,d)}(A)$, the closure of A, must intersect A.

Lemma 3.10

Let (X, d) be a metric space and let A be a non-empty subset of X. Suppose that V is a non-empty d-open subset of $\mathrm{Cl}_{(X,d)}(A)$. Then $A \cap V$ is non-empty.

Proof We suppose, by way of contradiction, that $A \cap V = \emptyset$. Now since V is d-open, its complement V^c is d-closed and hence $\mathrm{Cl}_{(X,d)}(A) \cap V^c$ is also d-closed. By our assumption, $A \subseteq V^c$ (and clearly $A \subseteq \mathrm{Cl}_{(x,d)}(A)$), so $A \subseteq \mathrm{Cl}_{(X,d)}(A) \cap V^c$. Therefore $\mathrm{Cl}_{(X,d)}(A) \cap V^c$ is a d-closed set containing A. Since, by definition, $\mathrm{Cl}_{(X,d)}(A)$ is the smallest d-closed set containing A, we have that $\mathrm{Cl}_{(X,d)}(A) \subseteq \mathrm{Cl}_{(X,d)}(A) \cap V^c$ and so $\mathrm{Cl}_{(X,d)}(A) \subseteq V^c$. That means that $V \subseteq (\mathrm{Cl}_{(X,d)}(A))^c$, which contradicts the assumption that V is a non-empty d-open subset of $\mathrm{Cl}_{(X,d)}(A)$. Thus $A \cap V$ is non-empty as required. ■

Before moving on, we pause for a moment to unpack exactly what we mean when we talk about a *connected subset* of X. Let (X, d) be a metric space, let $A \subseteq X$ and let d_A be the metric on A induced by the metric on X. When we talk about the connectedness of A, we want to consider whether we can find a closed and open subset E of A. The question then arises as to whether we want E to be d_A-closed and d_A-open, or d-closed and d-open. (For example, take $E = [1, 2)$, $A = (0, 2)$ and $X = \mathbb{R}$. Then E is closed when considered as a subset of A, but is not closed when considered a subset of $\mathbb{R}$.) To pre-empt any confusion, we make the following definition.

Definition 3.11 *Connected subspace*

Let (X, d) be a metric space with $A \subseteq X$. Then A is a **connected subspace** of X if (A, d_A) is connected when viewed as a metric space in its own right.

In other words, if (X, d) is a metric space and $A \subseteq X$, then we look for two subsets which are each simultaneously d_A-closed and d_A-open, when considered as subsets of A, in order to form a disconnection of A.

We now consider a metric space (X, d) and prove that if A is a d_A-connected subset of X, and $A \subseteq B \subseteq \mathrm{Cl}_{(X,d)}(A)$ then B is a d_B-connected subset of X. In other words, we want to prove that adding in any of the extra points needed to form the closure of a set cannot 'disconnect it'. We do this by contradiction. In the following proof we assume that B is d_B-disconnected, deduce there must be a (d_B, d_0)-continuous *onto* function $f\colon B \to \{0, 1\}$ and show that the restriction of this function to A is also onto (which is not possible since A is d_A-connected). We use Lemma 3.10 to show that the (d-open) preimage under f of *each* of the two points 0 and 1 must intersect A (and so $f|_A$ must be onto).

d_A and d_B are the metrics on A and B induced by the metric d on X.

Theorem 3.12

Let (X, d) be a metric space and let A be a d_A-connected subset of X. Write $\mathrm{Cl}(A)$ for $\mathrm{Cl}_{(X,d)}(A)$. Suppose that

$$A \subseteq B \subseteq \mathrm{Cl}(A).$$

Then B is d_B-connected.

In particular, $\mathrm{Cl}(A)$ is $d_{\mathrm{Cl}(A)}$-connected.

Proof The proof is by contradiction. Assume that B is d_B-disconnected. Then Theorem 3.9 implies that there is a (d_B, d_0)-continuous function f from B *onto* $\{0, 1\}$. Now the restriction of a continuous function between metric spaces is continuous (by Book D, Chapter 15, Theorem 1.6) and so $f|_A\colon A \to \{0, 1\}$, the restriction of f to A, is d_A-continuous.

Since $f\colon B \to \{0, 1\}$ is onto, we can find a point $b \in B$ such that $f(b) = 0$.

Consider $V = f^{-1}(\{0\})$, the preimage of $\{0\}$. Since $\{0\}$ is d_0-open in $\{0, 1\}$, we must have that V is d_B-open in B. Since V is d_B-open in B, by Theorem 2.1 we can find a d-open set V' with $V = V' \cap B$. Since V' is d-open, applying Theorem 2.1 a second time gives us a $d_{\mathrm{Cl}(A)}$-open set $U = V' \cap \mathrm{Cl}(A)$.

Since $B \subseteq \mathrm{Cl}(A)$, we have $b \in V = (V' \cap B) \subseteq (V' \cap \mathrm{Cl}(A)) = U$ and so U is a non-empty $d_{\mathrm{Cl}(A)}$-open subset of $\mathrm{Cl}(A)$. Here by Lemma 3.10 the set U has non-empty intersection with A (Figure 3.3), and we can find $a \in U \cap A$. Then $a \in U$, and so $f(a) = 0$.

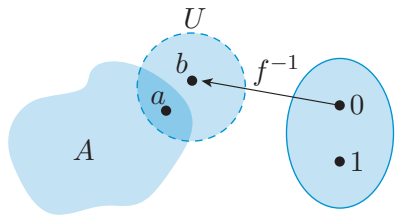

Figure 3.3

Since f is onto, there is also a point $b' \in B$ for which $f(b') = 1$.

Repeating the same argument, we conclude that there is a point $a' \in A$ for which $f(a') = 1$. Thus f, when restricted to A, is also onto. Since (A, d_A) is d_A-connected, this is a contradiction since no such continuous and onto function exists (by the remark following Theorem 3.9). Thus B is d_B-connected. ■

We have shown that the continuous image of a connected set is connected and that the closure of a connected set is connected. In the next section we think a little more about the nature of the pieces into which a set can fall.

3.3 Components

A disconnected space can often fall apart in different ways.

Exercise 3.10

Write down two different disconnections of $(\mathbb{Q}, d_{\mathbb{Q}})$.

For other spaces, such as $\left((0,1)\cup(1,2), d^{(1)}\right)$, there is only one obvious disconnection. In some sense, $\mathbb{Q}$ is more granular than $\left((0,1)\cup(1,2), d^{(1)}\right)$, which comes in two distinct 'lumps'. Notice that the metric is of fundamental importance here. Putting the discrete metric onto $(0,1)\cup(1,2)$ changes its nature entirely, since with the discrete metric, $(0,1)\cup(1,2)$ becomes at least as granular as $\mathbb{Q}$.

It is natural to ask whether there is an idea in which we can talk about the *largest* 'lumps' in a space. In other words, the largest connected subsets of a space.

Let (X,d) be a metric space and let $x \in X$. We know, from Exercise 3.3, that the single element set $\{x\}$ is connected. Therefore any given point x in X is always a member of at least one connected subset of X, namely $\{x\}$. In fact, it turns out that there is always a unique *largest* connected subset of X containing x.

We will pause briefly to clarify what we mean by 'largest' here. When we describe a set A as the *largest* connected subset of X containing x, we mean that if B is a connected subset of X, and $x \in B$, then $B \subseteq A$. In other words, A is a connected subset that contains any connected subset which contains x.

Definition 3.13 *Connected component*

Let (X,d) be a metric space and let $x \in X$.

The **d-component** $C_d(x)$ of x in X is the largest d-connected subset of X that contains x.

The existence of $C_d(x)$ is proved in Theorem 3.16.

Definition 3.14 *Totally disconnected*

X is **totally d-disconnected** if, for each point $x \in X$, the d-component of x is the set $\{x\}$.

Remark

When we do not wish to emphasise the metric, we say that X is totally disconnected.

It is not obvious that the component of a point always exists. To prove that it does, we first need a lemma in which we prove that under certain conditions the union of a family of connected subsets of a space X is connected. This is not true in general (for example, $X = (1, 2) \cup (3, 4)$ is disconnected). We need to make sure that the sets are 'stuck together' in a suitable way, with no set, or collection of sets, separated from the rest.

The following lemma states that the union, $A = \bigcup_{i \in I} A_i$, of a family of connected subsets is connected if the intersection of all the sets is non-empty. The proof is by contradiction. It may seem rather technical, but the idea is straightforward. We simply assume (for a contradiction) that the set A falls into two pieces, U and V, and show that this disconnection of A must cut at least one of the components, A_j, into two pieces $U_j = U \cap A_j$ and $V_j = V \cap A_j$ (contradicting the fact that A_j is connected). The technicalities come about in proving that all the relevant sets are non-empty.

Lemma 3.15

Let (X, d) be a metric space, and let $\{A_i : i \in I\}$ be a family of subsets of X and write d_{A_i} for the metric on A_i induced by d. Assume that the intersection of the family, $\bigcap_{i \in I} A_i$, is non-empty and that for each $i \in I$, A_i is d_{A_i}-connected. Then $A = \bigcup_{i \in I} A_i$ is d_A-connected (where d_A is the metric on A induced by the metric d on X).

Proof Note that all our sets are subsets of X and so all inherit the induced subspace metric from X. We want to prove that A is a d_A-connected subset of X. In other words, we want to show that (A, d_A) is d_A-connected, where d_A is the metric on A induced by the metric d on X. Therefore we work with respect to A and use the metric d_A.

Without loss of generality, we assume that the index set I has at least two elements, since otherwise there is nothing to prove. The proof is by contradiction.

Assume that A is d_A-disconnected, and so we can find two disjoint sets, U and V, each of which is simultaneously d_A-closed and d_A-open, with $A = U \cup V$ and $U \cap V = \emptyset$. Let a belong to the non-empty intersection of all the subsets $A_i = \bigcap_{i \in I} A_i$ (Figure 3.4). Then $a \in A$ must belong to either U or V. Without loss of generality, we may assume $a \in U$. So a is in the intersection of all of our sets A_i and is in U. Since V is non-empty, at least one of the sets A_i must meet V. Assume this is A_j. So the intersection $A_j \cap V$ is non-empty. Let $V_j = A_j \cap V$.
Since $a \in \bigcap_{i \in I} A_i$, we must have $a \in A_j$. Since $a \in A_j$ and $a \in U$, $A_j \cap U$ is also non-empty. Let $U_j = A_j \cap U$.

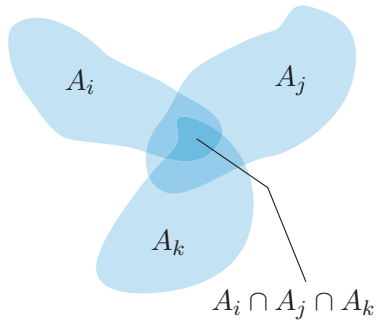

Figure 3.4

Consider the pair $\{U_j, V_j\}$. We have just shown them to be non-empty (by carefully selecting A_j so that V_j would be non-empty and by showing that $a \in U_j$). Now U is d_A-closed in A and so $U_j = U \cap A_j$ is d_{A_j}-closed in A_j. Similarly, V_j is d_{A_j}-closed in A_j.

They are also disjoint since U and V are disjoint and $U_j \subseteq U$ and $V_j \subseteq V$. Also their union is A_j since:

$$U_j \cup V_j = (A_j \cap U) \cup (A_j \cap V) = A_j \cap (U \cup V) = A_j \cap A = A_j.$$

So $A_j = U_j \cup V_j$ has been expressed as the disjoint union of two d_{A_j}-closed sets, and so is d_{A_j}-disconnected, which gives us our contradiction (as we assumed that each set A_j was d_{A_j}-connected). Therefore A is a d_A-connected subset of X. ■

We can now prove that the component of a point always exists and is well defined. Note that in the following proof the set $\mathcal{F}_x$ is a *set of sets*. An element of $\mathcal{F}_x$ is a connected *set* containing x. The component $C_d(x)$ is just one set, and is equal to the union of all the sets in $\mathcal{F}_x$.

Theorem 3.16

Let (X, d) be a metric space and let $x \in X$. Let

$$\mathcal{F}_x = \{A \subseteq X : x \in A, \text{ and } A \text{ is } d_A\text{-connected}\}.$$

Then $C_d(x)$, the d-component of x, exists and is given by

$$C_d(x) = \bigcup_{A \in \mathcal{F}_x} A.$$

Note $x \in C_d(x)$ since $\{x\} \in \mathcal{F}_x$.

Proof The single element set $\{x\}$ is in $\mathcal{F}_x$ and so $C_d(x)$ is not empty. Also since every set in $\mathcal{F}_x$ contains x (by definition), the intersection of all the sets in $\mathcal{F}_x$ is non-empty. So, by Lemma 3.15, $\bigcup_{A \in \mathcal{F}_x} A$ is connected.

If A is a connected subset of X that contains x, then A is one of the sets in $\mathcal{F}_x$, and so A must be part of the union of all of these sets. In other words, since A is an element of $\mathcal{F}_x$, then A is one of the sets in the union $\bigcup_{A \in \mathcal{F}_x} A$, which is $C_d(x)$. So $\bigcup_{A \in \mathcal{F}_x} A$ is the unique largest connected set that contains x. ■

Note that, if X is d-connected, then $C_d(x) = X$ for each $x \in X$.

Example 3.17

The components of $X = \left((0,1) \cup (2,3), d^{(1)}_{(0,1)}\right)$ are $(0,1)$ and $(2,3)$.

We prove that all open intervals are connected in Section 4.

Worked Exercise 3.18

Let $\mathbb{Q}$ have the induced metric $d^{(1)}_{\mathbb{Q}}$ from $\mathbb{R}$ with the Euclidean metric. Show that the component of each point $q \in \mathbb{Q}$ is $\{q\}$.

Solution

The proof is by contradiction. Let $q \in \mathbb{Q}$. Write d for the induced metric $d_{\mathbb{Q}}^{(1)}$ on $\mathbb{Q}$ induced by the Euclidean metric on $\mathbb{R}$. Assume, for a contradiction, that $C_d(q)$ consists of more than one point, so there is $p \neq q$ with $p \in C_d(q)$. Then, by an argument similar to that used in the solution to Worked Exercise 3.9 in Chapter 16, there is an irrational number x such that $\min\{p, q\} < x < \max\{p, q\}$. Set

$$U = \{r \in C_d(q) : r < x\} \quad \text{and} \quad V = \{r \in C_d(q) : r > x\}.$$

Then U and V are non-empty and disjoint, and their union is $C_d(q)$. Now $(-\infty, x) \cap \mathbb{Q}$ is d-open in $\mathbb{Q}$, and so $U = ((-\infty, x) \cap \mathbb{Q}) \cap C_d(q)$ is $d_{C_d(q)}$-open (where $d_{C_d(q)}$ is the metric on $C_d(q)$ induced by the metric d on $\mathbb{Q}$).

Similarly, V is $d_{C_d(q)}$-open for the induced metric on $C_d(q)$. Thus $\{U, V\}$ is a disconnection of $C_d(q)$. But this is impossible, since $C_d(q)$ is connected. This contradiction shows that there can be no point $p \neq q$ in $C_d(q)$. Thus, $C_d(q) = \{q\}$.

Since $\mathbb{Q}$ consists only of rational points, $x \notin C_d(q)$.

Do not be put off by the 'nested' induced metrics in the above example. We start with the Euclidean metric, $d^{(1)}$ on $\mathbb{R}$, and use this to induce a metric $d_{\mathbb{Q}}^{(1)}$ on $\mathbb{Q}$. We then take a subset $C_d(q)$ of $\mathbb{Q}$ and use $d_{\mathbb{Q}}^{(1)}$ to induce a metric on this subset of $\mathbb{Q}$ (which we write as $d_{C_d(q)}$).

Exercise 3.11

Describe the components of $X = \mathbb{R} - \mathbb{Z}$ with respect to the induced Euclidean metric. You may assume that $I = (a, b) \subseteq \mathbb{R}$ is $d_I^{(1)}$-connected. (We prove this in Section 4.)

Exercise 3.12

Describe the components of $(0, 1)$ with the discrete metric.

The components of a space are not (necessarily) the same as the sets forming a disconnection. For example, the sets $(-\infty, 1/2] \cap \mathbb{Q}$ and $(1/2, \infty) \cap \mathbb{Q}$ form a disconnection of $\mathbb{Q}$, as do the sets $(-\infty, \sqrt{3}) \cap \mathbb{Q}$ and $(\sqrt{3}, \infty) \cap \mathbb{Q}$, whereas the components of $\mathbb{Q}$ are the singleton sets $\{q\}$ for each $q \in \mathbb{Q}$. In particular, the components of a space need not be open although, as we see later, the components are always closed.

Worked Exercise 3.19

Show that the sets $\{q\}$, where $q \in \mathbb{Q}$, are not $d_{\mathbb{Q}}^{(1)}$-open in $\mathbb{Q}$ (where $d_{\mathbb{Q}}^{(1)}$ is the metric on $\mathbb{Q}$ induced by the Euclidean metric on $\mathbb{R}$).

Solution

Book D, Chapter 16, Theorem 2.9 states that $A \subseteq X$, where (X, d) is a metric space, is d-open if, and only if, we can find an $r > 0$ with $B_d(a, r) \subseteq A$ for all $a \in A$. Assume for a contradiction that $\{q\}$ is $d_{\mathbb{Q}}^{(1)}$-open in $\mathbb{Q}$. Clearly $q \in \{q\}$ so there must be an r with $B_{d_{\mathbb{Q}}^{(1)}}(q, r) \subseteq \{q\}$. In other words, we must have $B_{d_{\mathbb{Q}}^{(1)}}(q, r) = \{q\}$. Choose $n > \frac{1}{r}$. Then the rational number $q + \frac{1}{n} \in B_{d_{\mathbb{Q}}^{(1)}}(q, r)$, and so $B_d(q, r) \not\subseteq \{q\}$ and we have a contradiction.

Whilst a disconnected space can often be disconnected in many different ways, there is only one way to split it into its components. In fact, a set X is always the union of mutually disjoint components. We now prove this. In the following proof we define a relation, $\sim$, on a space X by letting $x \sim y$ if x and y are in the same component. We show that $\sim$ is an equivalence relation and so must partition the set into equivalence classes. These equivalence classes make up the individual components.

An equivalence relation always partitions a set into disjoint equivalence classes.

Theorem 3.20

Let (X, d) be a metric space. Let us write $x \sim y$ to mean that $y \in C_d(x)$. Then $\sim$ is an equivalence relation on X. Consequently, X is the union of its mutually disjoint d-components.

Proof In order to prove that $\sim$ is an equivalence relation we must show that $\sim$ is:

E1 Reflexive. For each $x \in X$, $x \sim x$.

E2 Symmetric. For each $x, y \in X$, if $x \sim y$ then $y \sim x$.

E3 Transitive. For each $x, y, z \in X$, if $x \sim y$ and $y \sim z$, then $x \sim z$.

Reflexive Since $x \in C_d(x)$, it follows immediately that $x \sim x$ for each $x \in X$.

Symmetric Suppose $x \sim y$ – that is, y is in $C_d(x)$, the component containing x (Figure 3.5). Now $C_d(x)$ is connected (by definition) and contains y and so must be contained in the largest connected subset of X containing y, which is $C_d(y)$ and so $x \in C_d(y)$. Therefore $\sim$ is symmetric.

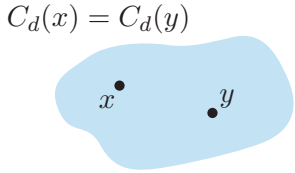

Figure 3.5 $x \sim y \Rightarrow y \sim x$

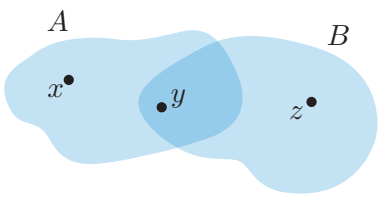

Figure 3.6 $x \sim y,\ y \sim z$

Transitive Suppose $x \sim y$ and $y \sim z$. Let A be a connected subset of X containing x and y, and B be a connected subset of X containing y and z (Figure 3.6). Note that $A \cap B$ contains y, and so is non-empty. Therefore, by Lemma 3.15, $A \cup B$ is connected. Therefore since x and z are in $A \cup B$, we have a connected set containing x and z and so it follows that $z \in C_d(x)$, and so $x \sim z$.

Hence $\sim$ is an equivalence relation, and so X can be written as the disjoint union of the equivalence classes defined by $\sim$. But, by definition of $\sim$, the equivalence class of a point x consists of all the points in $C_d(x)$ – namely $C_d(x)$ itself – and so each equivalence class is a component of X. ■

In particular, we conclude from Worked Exercise 3.18 that $\mathbb{Q}$ consists of countably many components, each of which is a singleton: $\mathbb{Q} = \bigcup_{q\in\mathbb{Q}}\{q\}$.

There is a useful immediate corollary to this theorem.

Corollary 3.21

Let (X, d) be a metric space. For each $y \in C_d(x)$, $C_d(y) = C_d(x)$. In other words, the component of each point in $C_d(x)$ is $C_d(x)$ itself.

3.4 The components are closed

From our investigation of $\mathbb{Q}$ in Worked Exercise 3.19 we know that the components of a space are not always open. In this section we show that they are always closed.

At the beginning of Section 3 we considered what happened when we moved the two intervals $(0, 1)$ and $(1, 2)$ closer and closer together. We described $X = (0, 1) \cup (1, 2)$ as disconnected and $(0, 1] \cup (1, 2) = (0, 2)$ as connected. We now consider the analogous situation in the plane.

Consider two discs of radius 1 in the plane, one centred at $(1, 1)$ and the second at $(3, 1)$ (see Figure 3.7).

First imagine that both the discs are open. In this case they do not meet at all, and the union is disconnected (as it is the disjoint union of the two open discs). If both are closed, then we expect their union to be connected (as in this case the discs meet in the point $(2, 1)$). So far, so good – but what if one is open and the other is closed? Does this third case describe a connected or disconnected subspace of the plane?

We can answer this question using the following lemma.

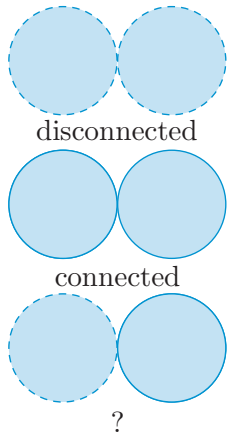

Figure 3.7

Lemma 3.22

Let (X, d) be a metric space.

(a) If $\{U, V\}$ is a disconnection of a metric space (X, d), then

$$\mathrm{Cl}_{(X,d)}(U) \cap V = \emptyset \text{ and } U \cap \mathrm{Cl}_{(X,d)}(V) = \emptyset.$$

(b) If there is a pair of non-empty sets U and V for which

$$U \cup V = X,\ \mathrm{Cl}_{(X,d)}(U) \cap V = \emptyset \text{ and } U \cap \mathrm{Cl}_{(X,d)}(V) = \emptyset,$$

then U and V are d-closed and form a disconnection of X.

In some texts, sets A_1 and A_2 for which $\mathrm{Cl}(A_i) \cap A_j = \emptyset$ $(i \neq j)$ are described as 'separated sets'.

Proof Part (a) follows from the definition of a disconnection since if $\{U, V\}$ is a disconnection of a metric space (X, d) then we know from Proposition 3.4 that both U and V are d-closed. Hence $\mathrm{Cl}_{(X,d)}(U) = U$ and $\mathrm{Cl}_{(X,d)}(V) = V$, and so $\mathrm{Cl}_{(X,d)}(U) \cap V = U \cap V = \emptyset$ and $U \cap \mathrm{Cl}_{(X,d)}(V) = U \cap V = \emptyset$.

For (b) suppose U and V are non-empty subsets of X, with

$$U \cup V = X, \ \mathrm{Cl}_{(X,d)}(U) \cap V = \emptyset \text{ and } U \cap \mathrm{Cl}_{(X,d)}(V) = \emptyset.$$

To show that U and V form a disconnection of X, it is sufficient to show that they are disjoint and d-closed (by Proposition 3.4). Since $U \subseteq \mathrm{Cl}_{(X,d)}(U)$, we must have $U \cap V = \emptyset$, so U and V are disjoint. Suppose that U is not d-closed. Since its complement in X is V, there are points in its d-closure that belong to V. But that contradicts the hypothesis that $\mathrm{Cl}_{(X,d)}(U)$ is disjoint from V; hence U is d-closed. In a similar way, we deduce that V is d-closed. Thus U and V form a disconnection of X. ■

Returning to our disc question – if X consists of two just touching discs in the plane, one $d^{(2)}$-open and one $d^{(2)}$-closed (where $d^{(2)}$ is the Euclidean metric on the plane), then the $d^{(2)}$-closure of the $d^{(2)}$-open disc meets the $d^{(2)}$-closed disc. Therefore the discs do not form a disconnection of X by Lemma 3.22. In fact, X is $d_X^{(2)}$-connected. Intuitively you may find it helpful to think of the open and closed discs as touching, but not meeting; whereas the two closed discs meet and the two open discs do not 'touch'.

We end this section by proving, as promised, that the components of a metric space (X, d) are d-closed.

This result, when combined with Theorem 3.20, tells us that a metric space can be split up into disjoint closed pieces, each of which is connected. These results are very useful when it comes to identifying the components of a set.

Lemma 3.23

Let (X, d) be a metric space and let $x \in X$. Then $C_d(x)$, the component of x in (X, d), is d-closed.

Proof This is Exercise 3.13. ■

Exercise 3.13

Prove Lemma 3.23.

Hint: you may find Theorem 3.12 useful.

3.5 Totally disconnected spaces

Totally disconnected spaces are spaces for which the components consist of single points (Definition 3.14). In this section we investigate some of these spaces. We have already met some such spaces. We saw in Exercise 3.12 that $\big((0,1), d_0\big)$ is totally d_0-disconnected and in Worked Exercise 3.18 we saw that $(\mathbb{Q}, d_{\mathbb{Q}}^{(1)})$ is totally disconnected.

Remarks

Although $\big((0,1), d_0\big)$ and $(\mathbb{Q}, d_{\mathbb{Q}}^{(1)})$ are both totally disconnected, they are very different as metric spaces. In $(\mathbb{Q}, d_{\mathbb{Q}}^{(1)})$ the distance between two points varies: $\frac{1}{2}$ is 'closer' to $\frac{24}{50}$ than it is to $\frac{1}{50}$; whereas the distance between $\frac{1}{2}$ and $\frac{24}{50}$ is the same as the distance between $\frac{1}{2}$ and $\frac{1}{50}$ in $\big((0,1), d_0\big)$. In $(\mathbb{Q}, d_{\mathbb{Q}}^{(1)})$ there is a very well defined sense of degree of 'closeness' which we do not have in any discrete space.

In addition, $\mathbb{Q}$ is countable and $(0,1)$ is uncountable.

We now investigate the middle-third Cantor set and show that this is another example of a totally disconnected space.

Worked Exercise 3.24

Recall that the middle-third Cantor set, K_C, is equal to $K_C = \bigcap_{n=1}^{\infty} K_n$ where:

$$K_0 = [0,1]$$
$$K_1 = \left[0, \tfrac{1}{3}\right] \cup \left[\tfrac{2}{3}, 1\right]$$
$$K_2 = \left[0, \tfrac{1}{9}\right] \cup \left[\tfrac{2}{9}, \tfrac{1}{3}\right] \cup \left[\tfrac{2}{3}, \tfrac{7}{9}\right] \cup \left[\tfrac{8}{9}, 1\right]$$
$$\vdots$$

with K_n formed by removing the open middle third of each interval in K_{n-1}.

Show that the middle-third Cantor set, K_C, with the metric induced by the Euclidean metric on $\mathbb{R}$, is totally disconnected.

See Book D, Chapter 16, Example 1.19.

Solution

We need to show that each component consists of a single element. Assume that $x \neq y$ are two elements belonging to the same component, C, of K_C. Without loss of generality, assume $x < y$. From the discussion of Example 1.19 in Chapter 16, there is a point $z \notin K_C$, and so certainly not in C, with $x < z < y$. Therefore the sets $(-\infty, z] \cap C$ and $[z, \infty) \cap C$ form a disconnection of C, which is a contradiction. Therefore our assumption that $x \neq y$ must be false.

4 Connectedness in Euclidean space

In this section, as promised, we consider which sets in $\mathbb{R}^n$ are connected. In what follows, all sets have metrics induced by the Euclidean metric unless otherwise stated.

Our first objective is to show that $\mathbb{R}$ is itself connected. This may seem intuitively obvious, but the proof is subtle. We need to prove that it is impossible to divide $\mathbb{R}$ into two sets, each of which is simultaneously open and closed. We know that we can do this easily with $\mathbb{Q}$, so the proof will need to use a property of $\mathbb{R}$ that $\mathbb{Q}$ does not possess. The required crucial property of $\mathbb{R}$ (that $\mathbb{Q}$ does not have) is the *least upper bound property*. We state it here for reference.

We often write $\sup A = M$.

If $A \subseteq \mathbb{R}$ is a non-empty set that is bounded above, then A has a **least upper bound**, or **supremum** – a real number M such that:

(a) $a \leq M$, for all $a \in A$

(b) if $M' < M$, then there is some $a \in A$ with $M' < a$.

We often write $\inf A = N$.

Similarly, if $A \subseteq \mathbb{R}$ is a non-empty set that is bounded below, then A has a **greatest lower bound**, or **infimum** – a real number N such that:

(a) $a \geq N$, for all $a \in A$

(b) if $N' > N$, then there is some $a \in A$ with $N' > a$.

Remark

Note that in the definition of supremum above, part (a) says that M is *an* upper bound for A and part (b) says that M is the *least* upper bound.

Remember that $\sqrt{2} \notin \mathbb{Q}$!

For example, the set $\{x \in \mathbb{R} : x^2 < 2\}$ has a least upper bound, namely $\sqrt{2}$ in $\mathbb{R}$; but $\{x \in \mathbb{Q} : x^2 < 2\}$ does not have a least upper bound in $\mathbb{Q}$.

Exercise 4.1

Use the least upper bound property to show that if a non-empty set $A \subseteq \mathbb{R}$ is bounded below, then the infimum exists.
Hint: consider the set $B = \{-a : a \in A\}$.

Remark

Note that a closed and bounded interval, $[a, b]$, contains its supremum and infimum.

We will use the least upper bound property to show that $\mathbb{R}$ is $d^{(1)}$-connected. First we prove a lemma.

Lemma 4.1

Let (X, d) be a metric space. A d-closed subset of X contains all of its d-closure points.

A point x is a closure point of a set A if every open ball centred on x meets the set A.

Proof Let A be a d-closed subset of X. Then $A = \mathrm{Cl}_{(X,d)}(A)$ by Theorem 3.3 of Chapter 16 and so, by Theorem 3.8 of Chapter 16, coincides with its set of closure points. ■

Theorem 4.2

$\mathbb{R}$ with the Euclidean metric is connected.

Proof The proof is by contradiction. We work with respect to the Euclidean metric, $d^{(1)}$. Suppose that $\mathbb{R}$ is disconnected, and let $\{U, V\}$ be a disconnection. Thus U and V are disjoint non-empty $d^{(1)}$-closed and $d^{(1)}$-open subsets of $\mathbb{R}$ with union $\mathbb{R}$.

Consider the pair of intervals $(-\infty, 0]$ and $[0, \infty)$. These are both $d^{(1)}$-closed and their union is $\mathbb{R}$, but they are not disjoint – they have the point 0 in common.

Since $U \cup V = \mathbb{R}$, the point 0 lies in one of the sets U and V. Without loss of generality, we can suppose that $0 \in V$. Since $U \neq \emptyset$, at least one of $(-\infty, 0] \cap U$ and $[0, \infty) \cap U$ is non-empty.

Case 1: $(-\infty, 0] \cap U \neq \emptyset$

Let $A = (-\infty, 0] \cap U$.

Since A is the intersection of two $d^{(1)}$-closed sets, it is $d^{(1)}$-closed. It is also bounded above by 0.

Since $0 \in V$ and $U \cap V = \emptyset$, $0 \notin U$, and so 0 cannot be in A. Hence

$$A = (-\infty, 0) \cap U.$$

In other words, removing the point 0 does not change A. As U is $d^{(1)}$-open, we now have A expressed as the intersection of two $d^{(1)}$-open sets and so A is $d^{(1)}$-open.

Hence A is $d^{(1)}$-closed and $d^{(1)}$-open, non-empty and bounded above (by 0). We now use these four properties of A to reach a contradiction to our assumption that $(-\infty, 0] \cap U \neq \emptyset$.

By the least upper bound property, A has a least upper bound – we denote this by M. We first show that M is a $d^{(1)}$-closure point of A.

Consider a $d^{(1)}$-open ball centred on M, $B(M, r)$ for an arbitrary $r > 0$. This contains $M - \frac{r}{2}$. Since $M - \frac{r}{2} < M$ we must be able to find an $a \in A$ with $M \geq a > M - \frac{r}{2}$ (because M is a *least* upper bound for A). Since $M \geq a > M - \frac{r}{2}$ we clearly have $a \in B(M, r)$.

Note that a may be equal to M.

Therefore $a \in A \cap B(M, r)$ and so $A \cap B(M, r)$ is non-empty. Therefore, since r was arbitrary, M is a $d^{(1)}$-closure point of A and so, by Lemma 4.1 M is an element of A.

Since A is also $d^{(1)}$-open, it contains a $d^{(1)}$-open ball centred at each of its points. In particular, it must contain a ball centred at M. So there is a number $r > 0$ such that $B(M, r) = (M - r, M + r) \subseteq A$.

But then A contains points y with $M < y < M + r$ – for example, $M + \frac{1}{2}r$ – and so M cannot be an upper bound for A.

We have arrived at a contradiction: so Case 1 does not apply.

Case 2: $[0, \infty) \cap U \neq \emptyset$

In this case, we let $A = [0, \infty) \cap U$, and use the greatest lower bound property to obtain a contradiction in a similar way: so Case 2 does not apply.

We conclude that $\mathbb{R}$ is connected. ■

Intuition suggests that every interval is connected. We ask you to prove this for the interval $(0, 1)$.

Exercise 4.2

Show that the interval $(0, 1)$ is a connected subset of $\mathbb{R}$ (with the Euclidean metric).

Exercise 4.3

Describe the components of $(0, 1)$ with the Euclidean metric.

Note that in Exercise 3.12 we showed that the components of $\big((0, 1), d_0\big)$ were the singleton sets $\{x\}$ for each $x \in (0, 1)$, whereas here we have shown that $\big((0, 1), d^{(1)}_{(0,1)}\big)$ has just one component. This highlights just how important the metric is. The space $\big((0, 1), d^{(1)}_{(0,1)}\big)$ comes in a single 'lump', whereas $\big((0, 1), d_0\big)$ falls into infinitely many tiny pieces!

It is now a simple step to show that any open interval is connected.

Proposition 4.3

Let $a < b$ be real numbers. The interval (a, b) is a connected subset of $\mathbb{R}$ (with the Euclidean metric).

Proof From the discussion at the beginning of Subsection 1.1, we know that any two open intervals of $\mathbb{R}$ are homeomorphic. As connectedness is preserved by homeomorphism, Exercise 4.2 shows that any open interval of $\mathbb{R}$ is connected. ■

Exercise 4.4

Let $a < b$ be real numbers. Show that the intervals $[a, b)$, $(a, b]$ and $[a, b]$ are connected subsets of $\mathbb{R}$ (with the Euclidean metric).

We have shown that all bounded intervals are connected. What about unbounded intervals such as (a, ∞)? The continuous image of a connected space is connected, so can we find a continuous map from $\mathbb{R}$ to (a, ∞)? It turns out that we can. The exponential map takes $\mathbb{R}$ to $(0, \infty)$ so we use this and rescale to get the map $\phi\colon \mathbb{R} \to (a, \infty)$ given by $\phi(x) = a + e^x$. This is continuous, showing that (a, ∞) is connected for each $a \in \mathbb{R}$. Likewise we can use the map $\psi(x) = a - e^x$ to show that $(-\infty, a)$ is connected. Theorem 3.12 now shows that $(-\infty, a]$ and $[a, \infty)$ are also connected.

Certainly $\emptyset$ and single-point subsets of $\mathbb{R}$ are connected and so we have shown that all the intervals in $\mathbb{R}$ are connected. In fact, it turns out that the *only* connected subsets of $\mathbb{R}$ are intervals.

Theorem 4.4

A subset A of $\mathbb{R}$ is connected for the Euclidean metric if, and only if, it is an interval.

Proof We have already shown that all intervals are connected. It remains to show that if A is connected, then it is an interval. Informally, if A is not an interval then it must have at least one 'hole' in it. In this proof by contradiction we use this hole to form a disconnection of A.

Let A be connected, and suppose that A is *not* an interval. Now the empty set and single-point sets are intervals, so A must contain at least two points. Since A is not an interval, there must be real numbers a, b, c with $a, b \in A$ and $a < c < b$ such that $c \notin A$ – otherwise A would be an interval. Let $\widetilde{E} = (-\infty, c]$ and $\widetilde{F} = [c, \infty)$. We now show that $E = \widetilde{E} \cap A$ and $F = \widetilde{F} \cap A$ form a disconnection of A.

The sets E and F are d_A-closed for the metric induced on A by the Euclidean metric on $\mathbb{R}$ since they are intersections with A of the $d^{(1)}$-closed sets $\widetilde{E}$ and $\widetilde{F}$. Neither is empty since $a \in \widetilde{E}$ and $b \in \widetilde{F}$. The intersection of E and F is

$$E \cap F = (\widetilde{E} \cap \widetilde{F}) \cap A = \emptyset,$$

and their union is

$$E \cup F = (\widetilde{E} \cup \widetilde{F}) \cap A = \mathbb{R} \cap A = A.$$

So we have two disjoint d_A-closed sets whose union is A. Therefore A is disconnected by Proposition 3.4.

This contradicts the fact that A is connected. We conclude that a connected subset of $\mathbb{R}$ must be an interval. ■

Exercise 4.5

Write down a disconnection of $\mathbb{R} - \{a\}$ (with the metric induced by the Euclidean metric on $\mathbb{R}$), for each $a \in \mathbb{R}$.

As connectedness is a topological invariant, we have the following corollary.

Corollary 4.5

$\mathbb{R} - \{a\}$ is not homeomorphic to $\mathbb{R}$ for the Euclidean metrics, for each $a \in \mathbb{R}$.

In fact, it follows from Theorem 4.4 that any subset of $\mathbb{R}$ that is not an interval must be disconnected.

The fact that intervals are connected can be used to show that many other sets are connected.

Exercise 4.6

Show that the unit circle $C = \{(x, y) : x^2 + y^2 = 1\}$ is connected.

It is now a short step to the promised generalisation of the Intermediate Value Theorem from Book D, Chapter 13.

Theorem 4.6 *Intermediate Value Theorem*

Let (X, d) be a connected metric space, and let $f : X \to \mathbb{R}$ be a $(d, d^{(1)})$-continuous function. Let a and b be points of X. Then f takes each value between $f(a)$ and $f(b)$.

Proof Since f is $(d, d^{(1)})$-continuous and X is connected, Lemma 3.6 implies that $f(X)$ is a connected subspace of $\mathbb{R}$, and so is an interval by Theorem 4.4. This interval contains $f(a)$ and $f(b)$, hence $f(X)$ contains each value between $f(a)$ and $f(b)$. ■

Remark

This remarkable theorem is seen to be an immediate consequence of three things: the domain is connected, the only connected subsets of $(\mathbb{R}, d^{(1)})$ are the intervals, and continuous images of connected sets are connected.

4.1 Homeomorphic subsets of $\mathbb{R}$: part 2

We can now conclude our investigation into which subsets of $\mathbb{R}$ are homeomorphic to each other. Recall that we had the following classes of intervals:

- the empty set $\emptyset$
- the singleton sets $\{a\} = [a, a]$
- the open intervals (a, b) for $a < b$, (a, ∞) and $(-\infty, a)$ for $a \in \mathbb{R}$, and $\mathbb{R}$
- the closed intervals $[a, b]$ for $a < b$
- the intervals of the form $[a, b)$ and $(a, b]$ for $a < b$ and $[a, \infty)$ and $(-\infty, a]$.

Intervals in the same class are homeomorphic. We now show that intervals in different classes are not homeomorphic.

Worked Exercise 4.7

Show that $A = (0, 1)$ is not homeomorphic to $B = (0, 1]$ (when both are considered as metric spaces with the metrics induced by the Euclidean metric on $\mathbb{R}$).

Solution

The proof is by contradiction. We know that A and B are both connected. Suppose that there is a homeomorphism $\phi\colon A \to B$. Then there must be a point in A which maps to the point $1 \in B$. Let $a \in A$ be this point (so $\phi(a) = 1$). Consider what happens when we take the point a away from $(0, 1)$, and also remove its image, the point 1, from $(0, 1]$. The restriction rule for homeomorphisms (Proposition 1.2), implies that $A - \{a\}$ is homeomorphic to $\phi\big(A - \{a\}\big) = B - \{1\} = A$. This is a contradiction, since $A - \{a\}$ is disconnected while A is connected. Thus A is not homeomorphic to B.

A disconnection is $\{(0, a), (a, 1)\}$.

Since $(0, 1)$ is homeomorphic to (a, b) and $(0, 1]$ is homeomorphic to $(a, b]$ and to $[c, d)$ (where all spaces have metrics induced by the Euclidean metric on $\mathbb{R}$), it follows from Worked Exercise 4.7 that, in general, an open interval and a half-open interval are not homeomorphic (when both are considered as metric spaces with the metric induced by the Euclidean metric on $\mathbb{R}$).

Exercise 4.7

Show that $A = (0, 1]$ is not homeomorphic to $B = [0, 1]$ (when both are considered as metric spaces with the metric induced by the Euclidean metric on $\mathbb{R}$).

It follows that, in general, $(a, b]$ and $[a, b]$ are not homeomorphic (when both are considered as metric spaces with the metric induced by the Euclidean metric on $\mathbb{R}$).

Remarks

1. Intuitively we would expect different types of interval to be different since the endpoints 'look' different. Imagine that you are a one-dimensional creature living inside the interval $[0, 1)$. The left-hand side would have a hard edge which you could bump right into. However, you could never reach the right-hand boundary. The same is true of the left- and right-hand edges of $[0, \infty)$, so it makes intuitive sense for the two intervals to be homeomorphic.
2. Whilst the intervals $[a, b)$ and $[a, \infty)$ share many topological properties, they do not, of course, share all properties: for instance, $[a, b)$ is bounded whilst $[a, \infty)$ is not. Homeomorphism allows us to stretch space as much as we like (and so, as we have seen, a bounded interval can be homeomorphic to an unbounded one).

4.2 Products of connected spaces

After this successful investigation of the connectedness of the real line, we turn our attention to the connectedness (or otherwise) of Euclidean space, $\mathbb{R}^n$ with $n > 1$.

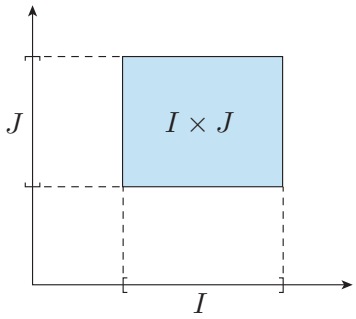

Figure 4.1

The fact that an n-fold product of connected spaces is also connected then follows by induction.

We saw in Book D, Chapters 13 and 14 that many of the properties of $\mathbb{R}^n$ are immediate consequences of the fact that $\mathbb{R}^n$ is formed by repeatedly taking products of $\mathbb{R}$. In particular, the Euclidean metric on $\mathbb{R}^2$ is a product metric of the Euclidean metric on $\mathbb{R}$ with itself.

It seems plausible that connectedness should be preserved when we take products – for example, the product of an interval with an interval gives a rectangle in the plane (Figure 4.1), a set that we expect to be connected.

In order to show that $\mathbb{R}^2$ (or a rectangle) is connected, we prove a more general result: that the product of any two connected metric spaces is itself a connected metric space.

Theorem 4.8

Let (X, d_1) and (Y, d_2) be non-empty metric spaces. The product space $(X \times Y, e)$, where e is a product metric, is e-connected if, and only if, (X, d_1) and (Y, d_2) are both connected (for their respective metrics).

Proof We prove that if X and Y are connected for their respective metrics, then $X \times Y$ is e-connected. We leave the converse as an exercise for you (Exercise 4.8).

Suppose that (X, d_1) and (Y, d_2) are connected (for their respective metrics). We must show that $(X \times Y, e)$ is e-connected. To do this, it is enough to show that, for a given $z \in X \times Y$, the component of z is the whole space $X \times Y$.

So suppose that $z = (x, y) \in X \times Y$, and let $C_e(z)$ denote the component of z in $X \times Y$. In order to show that $C_e(z) = X \times Y$, we must show that if $z' \in X \times Y$ is any other point in $X \times Y$ then $z' \in C_e(z)$.

So suppose that $z' = (x', y') \in X \times Y$. We show that $z' \in C_e(z)$ by finding a connected subset A of $X \times Y$ that contains z and z'. As the component containing (x, y), by definition, contains all other connected subsets of $X \times Y$, then we are done. We define $A = (X \times \{y\}) \cup (\{x'\} \times Y)$, as illustrated in Figure 4.2.

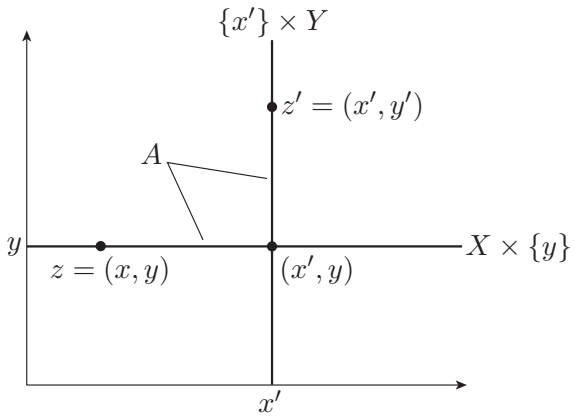

Figure 4.2 $A = (X \times \{y\}) \cup (\{x'\} \times Y)$

Clearly, both z and z' are elements of A, so it remains only to show that A is e-connected (as a subset of $X \times Y$).

Now $X \times \{y\}$ (as a subspace) is a $(d_1, e_{X \times \{y\}})$-continuous image of X and $\{x'\} \times Y$ is a $(d_2, e_{\{x'\} \times Y})$-continuous image of Y. Thus, by Lemma 3.6, both $X \times \{y\}$ and $\{x'\} \times Y$ are connected.

We do not include the proofs of these straightforward but technical results.

However, the point (x', y) is in both $X \times \{y\}$ and $\{x'\} \times Y$. Thus, by Lemma 3.15, A is connected. ■

Exercise 4.8

Prove the second part of this theorem, that if (X, d_X) and (Y, d_Y) are non-empty metric spaces and $(X \times Y, e)$ is e-connected, where e is a product metric, then (X, d_X) and (Y, d_Y) are connected (for their respective metrics). *Hint*: if (X, d_X) and (Y, d_Y) are metric spaces, and e is a product metric for $X \times Y$, then the projection function $p_X : X \times Y \to X$ is (e, d_X)-continuous and the projection function $p_Y : X \times Y \to Y$ is (e, d_Y)-continuous.

See Book D, Chapter 15, Theorem 3.11.

We have two immediate corollaries.

Corollary 4.9

Euclidean space, $(\mathbb{R}^n, d^{(n)})$, is connected.

Corollary 4.10

The n-dimensional 'rectangles' $I_1 \times \cdots \times I_n$ are connected for the Euclidean metric, where I_i is an interval in $\mathbb{R}$.

Exercise 4.9

Let $f\colon \mathbb{R}^2 \to \mathbb{R}$ be the continuous polynomial defined by

$$f(x_1, x_2) = 5x_1^5 + 5x_1^3x_2^2 - x_2^5.$$

Show that there is a point in the plane at which f attains the value 42.

We now know enough to deduce that many other spaces are connected. The following discussion is informal but should give you an idea of how the above results can be applied.

We know that the unit circle S is a connected subset of the plane from Exercise 4.6. What happens if we take a product of a circle with itself? Imagine taking a flexible cylindrical pipe (such as a section of particularly bendy hosepipe).

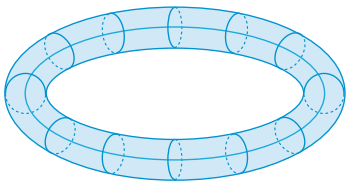

Figure 4.3

Now bend this around and imagine joining the two ends together seamlessly. This would produce a circular pipe (in the sense that placing the circular pipe onto sand would leave a circular imprint). The shape of this circular pipe would be similar to that of a child's rubber bathing ring (Figure 4.3). If you cut through this circular pipe at any point then you would see the cross-section of the original hose – in other words, a circle. So we can think of this surface as being the product of two circles. Mathematically this surface is a torus. Therefore we can now deduce that the torus is connected.

The punctured circle is homeomorphic to the Euclidean line $\mathbb{R}$, by means of stereographic projection (Figure 4.4(a)). Similarly, the punctured sphere is homeomorphic to the Euclidean plane (Figure 4.4(b)). Since the Euclidean plane is connected, so is the punctured sphere. The puncture is repaired by the operation of *closure*, and since closure preserves connection, by Theorem 3.12, the sphere is connected.

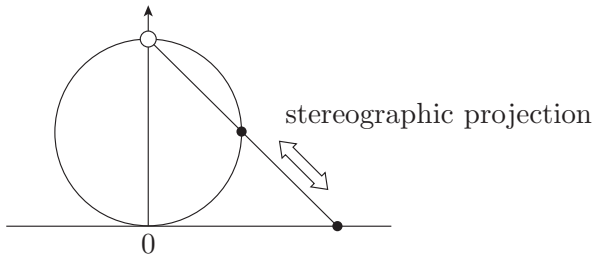

(a)

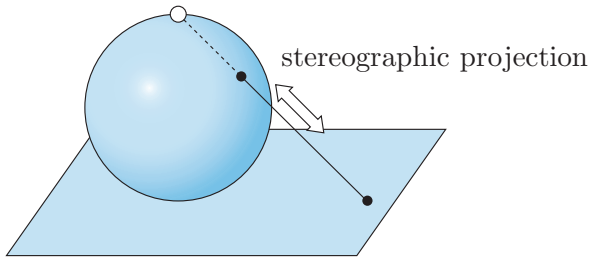

(b)

Figure 4.4

5 Path-connected spaces

Basing our notion of connectedness on whether or not a space falls into pieces seems very natural if we consider looking at a space from the 'outside'. However, if we lived *inside* the space, then we might come up with an alternative concept. We may decide that a space is connected if we can get from any one point to any other point without leaving our space. This alternative description of being connected leads to the idea of *path-connectedness*. Intuitively we say that a space is *path-connected* if it is possible to find a 'path' between any two of its points (Figure 5.1).

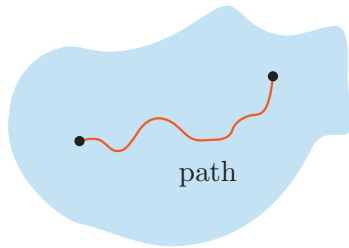

Figure 5.1

We cannot proceed any further until we define what we mean by a *path* in a metric space. The concept of a path is of wide interest. If you have studied a course in complex analysis (contour integration), fluid mechanics (vortices), electromagnetism (the integral laws relating currents and magnetic fields), or differential geometry and its applications to general relativity, you have met path (or line) integrals, and hence paths. Think for a moment about how you could define a path ....

... You may have thought about something that you can draw without taking your pen off the page. This is a useful way of thinking about a path in $\mathbb{R}^2$, but not so helpful when considering a path in, say, a space of functions. Writing definitions is often tricky and mathematicians spend a fair amount of their time trying to correctly define their terms. Indeed, getting the definitions right can be the hardest part of proving a result.

5.1 Defining a path

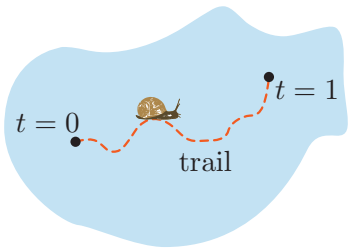

Figure 5.2

What should we mean by a path? We certainly want our idea of a mathematical path to include the intuitive idea of a path on the plane, so we start there. Imagine a snail crawling around in $\mathbb{R}^2$ (Figure 5.2). As the snail moves, it leaves a trail behind it on the surface. Suppose the motion takes place from time $t = 0$ to $t = 1$, after which we take the snail off the surface and examine the trail left behind.

What will not be apparent is how the trail was made: did the snail speed up, or slow down? Did it stop for a rest at any point? Did it reverse back along a portion of the trail, then reverse again and go on its way? There are infinitely many ways that a given trail can be laid down, so the concept of a trail is inadequate for mathematical purposes.

A more informative way to describe the snail's progress is to record where it is at each instant in time – that is, instead of focusing on the trail, we consider the *function* that records the snail's position as a function of time. Mathematically, this is a function $p\colon [0,1] \to X$, where X is the surface upon which the snail moves. The interpretation is that the snail is at the point $p(t) \in X$ at time $t \in [0,1]$: in particular, it starts at $p(0)$ and finishes at $p(1)$. In this language, the trail is the image set $p([0,1]) \subseteq X$.

The only condition we demand of a path is that it be a $(d^{(1)}_{[0,1]}, d)$-*continuous function*: the snail may speed up, slow down, stop, reverse, change direction, as often and as abruptly as it pleases, but it must not 'break' its motion by, for example, jumping along at any point.

The snail is not a grasshopper!

The trail may be extremely complicated and hard to visualise. It is known, for example, that there are paths on the plane whose trails completely fill the unit square! However, the advantage of thinking of a path as a function is that the fact that our path must not have any breaks in it can be encapsulated in the idea that the function must be continuous. If we focused on the image set (the trail) then we could possibly come up with a workable definition for the plane, but it would be harder to generalise our definition to paths in other spaces. With this definition, there is no reason as to why we should confine ourselves to paths on surfaces in $\mathbb{R}^2$. All that we require is continuity, and so everything we have said so far holds if the codomain is an arbitrary metric space (except, of course, for the bits about the snail).

The Peano curve is such an example (see Figure 5.3).

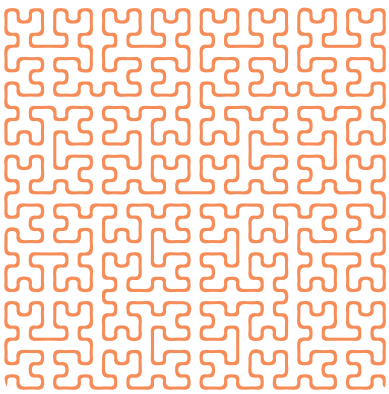

Figure 5.3

Definition 5.1 *Path*

Let (X, d) be a metric space and let $[0,1]$ have the induced Euclidean metric $d^{(1)}_{[0,1]}$.

A function $p\colon [0,1] \to X$ is a **path** in X if p is a $(d^{(1)}_{[0,1]}, d)$-continuous function (Figure 5.4). A path p **goes from** a **to** b in X if p is a path in X and if there are t_x, $t_y \in [0,1]$ with $t_x \leq t_y$, $p(t_x) = a$ and $p(t_y) = b$. If the order in which the points are reached is unimportant, then we say that p **joins** a and b.

The **initial point** of a path p is the point $p(0)$ in X and the **final point** of p is the point $p(1)$.

The path is a **closed path** if the initial and final points of a path are the same (that is, $p(0) = p(1)$).

The **trail** of a path is its image set $p([0,1])$.

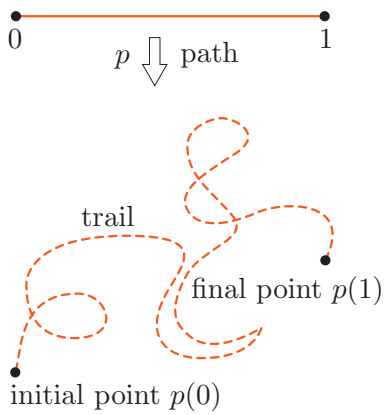

Figure 5.4

Remarks

1. The closed nature of the domain ensures that our path has a definite starting point and ending point.
2. Starting with a connected domain enables us to use Lemma 3.6 to guarantee that the trail $p([0,1])$ is a connected subset of X.
3. The meaning of 'closed' in 'closed path' is different from its meaning in 'closed set'. This unfortunate coincidence of terminology is standard. It is usually clear from the context which meaning of 'closed' is intended.
4. For each point x in a metric space (X, d), there is always a path joining a point x to itself – namely, the constant function $p\colon [0,1] \to X$ given by $p(t) = x$.

Exercise 5.1

Let (X, d) be a metric space. Prove part 4 of the remarks above. In other words, prove that the constant function $p\colon [0,1] \to X$ given by $p(t) = x$ is $(d^{(1)}_{[0,1]}, d)$-continuous.

Worked Exercise 5.2

Let $\mathbb{R}^2$ have the Euclidean metric. Find a path with initial point $(0,0)$ and final point $(1,1)$.

Solution

The simplest example of such a path is probably the function $p\colon [0,1] \to \mathbb{R}^2$ given by $p(t) = (t,t)$, illustrated in Figure 5.5. We need to verify that this p is $(d^{(1)}_{[0,1]}, d^{(2)})$-continuous. In order to do so, we note that $p_1 \circ p$ and $p_2 \circ p$ are both $(d^{(1)}_{[0,1]}, d^{(1)})$-continuous functions from $[0,1]$ to $\mathbb{R}$, where p_1 and p_2 denote the usual projection functions from $\mathbb{R}^2$ to $\mathbb{R}$. Hence, by Theorem 3.12 of Chapter 15, the function p is $(d^{(1)}_{[0,1]}, d^{(2)})$-continuous. Since $p(0) = (0,0)$ and $p(1) = (1,1)$, the map p gives a path with initial point $(0,0)$ and final point $(1,1)$.

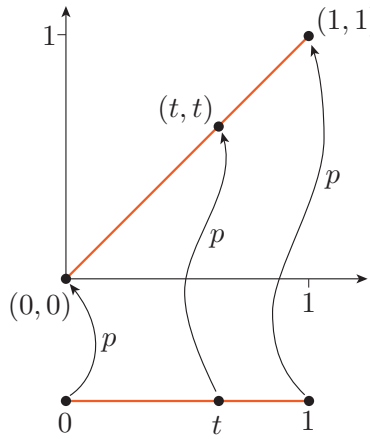

Figure 5.5

Exercise 5.2

Let $X = [0,1] \times [0,1]$ with the induced Euclidean metric. Show that there is a path joining $(0,0)$ and $(1, \frac{1}{2})$ in X.

We next show that paths can be joined. This may seem obvious, but we need to ensure that we can join two paths together and still end up with a continuous map from $[0,1]$ to X, since our definition of path requires this. We first need a lemma.

Lemma 5.3

Suppose we have maps $f\colon [a,b] \to \mathbb{R}$ and $g\colon [b,c] \to \mathbb{R}$, with $f(b) = g(b)$ and where f and g are continuous with respect to the metrics induced by the Euclidean metric on $\mathbb{R}$.
Then $h\colon [a,c] \to \mathbb{R}$ defined by:

$$h(t) = \begin{cases} f(t), & \text{for } a \leq t \leq b, \\ g(t), & \text{for } b \leq t \leq c, \end{cases}$$

is also continuous with respect to the metrics induced by the Euclidean metric on $\mathbb{R}$.

Proof

Throughout this proof, we work with respect to the metrics induced by the Euclidean metric on $\mathbb{R}$. We write d for $d^{(1)}$. By Theorem 7.1 of Chapter 16, the function $h\colon [a,c] \to \mathbb{R}$ will be $(d_{[a,c]}, d)$-continuous if, and only if, $h^{-1}(F)$ is a $d_{[a,c]}$-closed set in $[a,c]$ whenever F is a d-closed set in $\mathbb{R}$.

First we note the following: a set $A \subseteq [r,s]$ is a $d_{[r,s]}$-closed subset of $[r,s]$ if, and only if, we can find a d-closed subset of $\mathbb{R}$, say E, with $A = [r,s] \cap E$. Now $[r,s]$ is d-closed in $\mathbb{R}$ and so A, being the intersection of two d-closed sets, is d-closed. So A is both a d-closed subset of $\mathbb{R}$ and a $d_{[r,s]}$-closed subset of $[r,s]$.

Let F be a d-closed set in $\mathbb{R}$. Then since $f(b) = g(b)$ we have:

$$\begin{aligned} h^{-1}(F) &= \{x \in [a,c] : h(x) \in F\} \\ &= \{x \in [a,b] : f(x) \in F\} \cup \{x \in [b,c] : g(x) \in F\} \\ &= f^{-1}(F) \cup g^{-1}(F). \end{aligned}$$

But f is $(d_{[a,b]}, d)$-continuous and g is $(d_{[b,c]}, d)$-continuous, and so by Theorem 7.1 of Chapter 16, $f^{-1}(F)$ is a $d_{[a,b]}$-closed subset of $[a,b]$ and $g^{-1}(F)$ is a $d_{[b,c]}$-closed subset of $[b,c]$. Therefore, by the note near the start of this proof, $f^{-1}(F)$ and $g^{-1}(F)$ are both d-closed subsets of $\mathbb{R}$ and so their union $f^{-1}(F) \cup g^{-1}(F)$ is a d-closed subset of $\mathbb{R}$. Hence, by Theorem 3.12 of Chapter 16, $f^{-1}(F) \cup g^{-1}(F) = f^{-1}(F) \cup g^{-1}(F) \cap [a,c]$ is a $d_{[a,c]}$-closed subset of $[a,c]$, and h is $(d_{[a,c]}, d)$-continuous. ■

Lemma 5.4

Let (X, d) be a metric space and let $a, b, c \in X$. Suppose that there is a path p in X with initial point a and final point b, and a path q in X with initial point b and final point c. Then there is a path that joins a and c in X; see Figure 5.6.

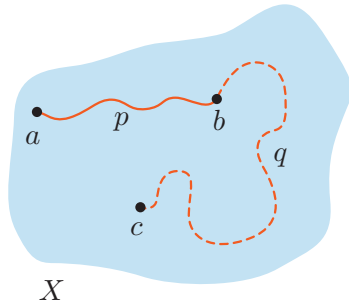

Figure 5.6

Proof In this proof, all metrics are the metrics induced by the Euclidean metric on $\mathbb{R}$.

The basic idea here is simple: we define a path joining a and c by doing p followed by q. There is a problem, however: both p and q are defined for $t \in [0, 1]$, and what we would like is to trace the whole of p for $t \in [0, \frac{1}{2}]$ and then trace the whole of q for $t \in [\frac{1}{2}, 1]$, as illustrated in Figure 5.7. The way to do this is to trace out both p and q at double their original 'speed'. Formally, we do this as follows.

Define $r\colon [0,1] \to X$ by

$$r(t) = \begin{cases} p(2t), & \text{for } 0 \leq t \leq \frac{1}{2}, \\ q(2t-1), & \text{for } \frac{1}{2} \leq t \leq 1. \end{cases}$$

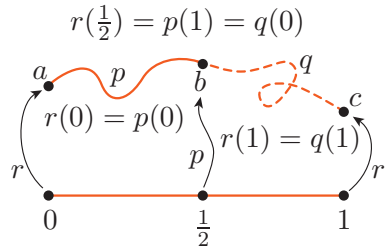

Figure 5.7

Certainly $r(0) = p(0) = a$ and $r(1) = q(1) = c$ and so it remains only to show that r is $(d^{(1)}_{[0,1]}, d)$-continuous. The map taking t to $p(2t)$ is the composition of p with the continuous map $f\colon t \mapsto 2t$, and the map taking t to $q(2t-1)$ is the composition of q with the continuous map $g\colon t \mapsto 2t-1$. Therefore both are continuous (by the Composition Rule for continuous functions in Book D, Chapter 13). Also $r(\frac{1}{2}) = p(1) = q(0) = b$. Therefore r is continuous by Lemma 5.3. ■

Now that we have the notion of a path, we can define what we mean for a metric space to be path-connected.

Definition 5.5 *Path-connected*

Let (X, d) be a metric space.

The space (X, d) is **path-connected** if, for each a and b in X, there is a path in X that joins a and b.

A set $A \subseteq X$ is **path-connected** if (A, d_A) is path-connected.

Recall that the induced metric on A is denoted by d_A.

Remarks

1. If $X = \emptyset$, then X is certainly path-connected, since there are no points to join.
2. If $X = \{x\}$ has only one point, the constant function $f\colon [0,1] \to X$ given by $f(t) = x$ is a path, and so X is path-connected.
3. When we wish to emphasise the metric, we use the term *d-path-connected*.

In the next exercise, we ask you to show that paths can be traced backwards.

Exercise 5.3

Let (X, d) be a metric space and let $a, b \in X$. Suppose that $p\colon [0,1] \to X$ is a path in X with initial point a and final point b. Show that there is a path q in X with initial point b and final point a.

5.2 Examples of path-connected spaces

We now look at some examples.

Worked Exercise 5.6

Show that $\mathbb{R}$ with the Euclidean metric is path-connected.

Solution

In this solution, all metrics are the metrics induced by the Euclidean metric on $\mathbb{R}$.

Let a, $b \in \mathbb{R}$; we must find a path joining a and b in $\mathbb{R}$. We already know how to define such a path if $a = b$, so suppose that $a \neq b$. Without loss of generality, we shall assume that $a < b$. There are many paths that we can choose. The standard one is the straight-line path $P\colon [0,1] \to \mathbb{R}$ defined by

$$P(t) = (1-t)a + tb,$$

illustrated in Figure 5.8. As P is simply the restriction of a linear function to the interval it is clearly continuous. Also we have $P(0) = a$ and $P(1) = b$. Hence P is a path in $\mathbb{R}$ from a to b and $\mathbb{R}$ is path-connected.

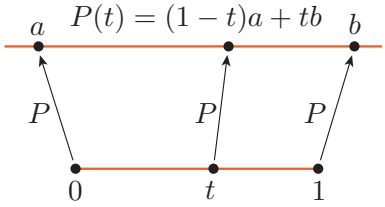

Figure 5.8

Exercise 5.4

Show that $[a, b]$ with the induced Euclidean metric is path-connected.

Remark

The result of Exercise 5.4 enables us to deduce that every interval I in $\mathbb{R}$ is path-connected. We know that the empty interval is path-connected. If I is not empty, for each a and b in I $(a \leq b)$ we know that $[a, b]$ is path-connected and hence there is a path from a to b in $[a, b] \subseteq I$. Thus I is path-connected.

The result in Worked Exercise 5.6 generalises to $\mathbb{R}^n$ (for any $n \in \mathbb{N}$).

Proposition 5.7

The metric space $\mathbb{R}^n$ with the Euclidean metric is path-connected.

Proof For $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$ we define $P\colon [0,1] \to \mathbb{R}^n$ by $P(t) = (1-t)\mathbf{a} + t\mathbf{b}$.

We must prove that P is a path from $\mathbf{a}$ to $\mathbf{b}$.

Continuity: let $p_j\colon \mathbb{R}^n \to \mathbb{R}$ be the jth projection function, defined by $p_j(\mathbf{a}) = a_j$. Then $(p_j \circ P)(t) = (1-t)a_j + tb_j$, so $p_j \circ P$ is linear and so clearly continuous. Since $p_j \circ P$ is continuous for $j = 1, 2, \ldots, n$, it follows from the generalisation of Theorem 3.12 of Chapter 15 that P is continuous.

Endpoints: by inspection, $P(0) = \mathbf{a}$ and $P(1) = \mathbf{b}$. Thus P is a path from $\mathbf{a}$ to $\mathbf{b}$ and $\mathbb{R}^n$ is path-connected. ■

Exercise 5.5

In the proof of Proposition 5.7 we have suppressed the metrics and written 'continuous' rather than (d, e)-continuous (for appropriate metrics d and e). Insert the appropriate metrics for each use of the word 'continuous' in this proof.

We have seen that connectedness is preserved under continuous mappings, and hence is a topological invariant. The same is true for path-connectedness.

Theorem 5.8

Let (X, d_X) and (Y, d_Y) be metric spaces, let X be path-connected, and let $f\colon X \to Y$ be (d_X, d_Y)-continuous. Then $f(X)$ is path-connected.

Proof Let $a, b \in f(X)$. By the definition of $f(X)$, there exist $c, d \in X$ such that $f(c) = a$ and $f(d) = b$. Since X is path-connected, there exists a path $p\colon [0,1] \to X$ from c to d. Since f is (d_X, d_Y)-continuous and p is $(d^{(1)}_{[0,1]}, d_X)$-continuous, $f \circ p\colon [0,1] \to f(X)$ is a path from $f(c) = a$ to $f(d) = b$ in $f(X)$ (with the induced metric inherited from (Y, d_Y)). Thus $f(X)$ is path-connected. ■

Corollary 5.9

Path-connectedness is a topological invariant.

We can use the methods of the proof of Proposition 5.7 to show that every $d^{(n)}$-open ball in $\mathbb{R}^n$ is path-connected for the Euclidean metric. The first step is to show that the unit open ball centred at $\mathbf{0}$ in $\mathbb{R}^n$ is path-connected for the Euclidean metric. We then use the fact that path-connectedness is a topological invariant to show that this applies to every $d^{(n)}$-open ball.

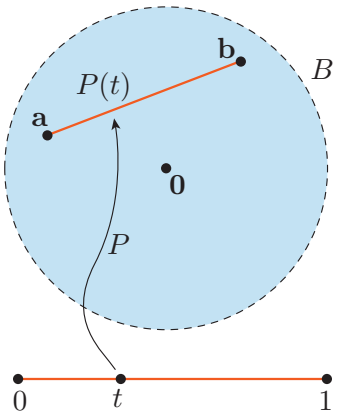

Figure 5.9

Worked Exercise 5.10

Show that the unit $d^{(n)}$-open ball B in $\mathbb{R}^n$ with the induced Euclidean metric is path-connected.

Solution

Let $\mathbf{a}, \mathbf{b} \in B$. We know from the proof of Proposition 5.7 that $P(t) = (1-t)\mathbf{a} + t\mathbf{b}$ defines a path from $\mathbf{a}$ to $\mathbf{b}$ in $\mathbb{R}^n$ (Figure 5.9). Therefore we can deduce that P defines a path in B if we can show that $P(t)$ belongs to B for all $t \in [0,1]$ – that is, we need to show that $d^{(n)}\big(P(t), \mathbf{0}\big) < 1$. We have:

$$\begin{aligned} d^{(n)}\big(P(t), \mathbf{0}\big) &= d^{(n)}\big((1-t)\mathbf{a} + t\mathbf{b}, \mathbf{0}\big) \quad \text{(definition of } P\text{)} \\ &\leq d^{(n)}\big((1-t)\mathbf{a} + t\mathbf{b}, t\mathbf{b}\big) + d^{(n)}\big(t\mathbf{b}, \mathbf{0}\big) \quad \text{(Triangle Inequality)} \\ &= d^{(n)}\big((1-t)\mathbf{a}, \mathbf{0}\big) + d^{(n)}\big(t\mathbf{b}, \mathbf{0}\big) \quad \text{(translating by } t\mathbf{b}\text{)} \\ &= (1-t)d^{(n)}\big(\mathbf{a}, \mathbf{0}\big) + td^{(n)}\big(\mathbf{b}, \mathbf{0}\big) \quad \text{(factorising)} \\ &< (1-t) + t = 1. \quad (\mathbf{a}, \mathbf{b} \in B) \end{aligned}$$

Therefore B is path-connected.

A proof almost identical to that in Worked Exercise 5.10 allows us to show that the unit $d^{(n)}$-closed ball in $\mathbb{R}^n$ is path-connected, and thus to deduce that every $d^{(n)}$-closed ball in $\mathbb{R}^n$ is path-connected.

Theorem 5.11

Every $d^{(n)}$-open ball in $\mathbb{R}^n$ is path-connected. Every $d^{(n)}$-closed ball in $\mathbb{R}^n$ is path-connected.

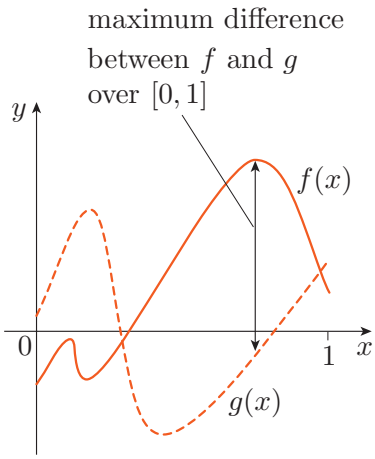

Figure 5.10

So far, we have just considered path-connectedness for the Euclidean metric. We now consider a different metric space.

Recall that $C[0,1]$ is the set of continuous functions from $[0,1]$ to $\mathbb{R}$ and that the max metric on $C[0,1]$ is defined by

$$d_{\max}(f,g) = \max\{|g(x) - f(x)| : x \in [0,1]\}.$$

It may seem a little strange to think of a path in a space of functions but there is no mathematical reason not to. Our metric, $d_{\max}$, gives us a sense of the distance between two functions, where the distance between f and g is the maximum vertical distance between them (as shown in Figure 5.10).

One function is 'close' to another if the maximum vertical distance between the two graphs of the functions is 'small'. Intuitively, we can think of a trail, $p\big([0,1]\big)$, from the function f to the function g as being a series of functions, $p(t) \in C[0,1]$, such that drawing the graphs of each would give us a way of deforming the graph of f into the graph of g. This is not a rigorous definition, but may be a helpful way of visualising a trail in a function space.

Worked Exercise 5.12

Show that $\big(C[0,1], d_{\max}\big)$ is path-connected.

Solution

Let f and g be functions in $C[0,1]$. Consider the function $P\colon [0,1] \to C[0,1]$ given by

$$P(t) = (1-t)f + tg,$$

so that, for $x \in [0,1]$ and $t \in [0,1]$,

$$P(t)(x) = (1-t)f(x) + tg(x).$$

We ask you to show that P is $(d^{(1)}_{[0,1]}, d_{\max})$-continuous in Exercise 5.6. Note that $P(0) = f$ and $P(1) = g$. We deduce that P is a path from f to g, and $C[0,1]$ is path-connected.

Exercise 5.6

By bounding $d_{\max}\big(P(t), P(s)\big)$ in terms of $|t-s|$ and $d_{\max}(f,g)$, prove that P is $(d^{(1)}_{[0,1]}, d_{\max})$-continuous.

You may have noticed that all the examples of path-connected spaces that you have met so far are also connected. This is no accident, as the following theorem shows.

Theorem 5.13

Let (X,d) be a metric space. If X is d-path-connected, then X is d-connected.

Proof We need to show that if we can find a path between any two points in our space then we cannot form a disconnection. If X is empty, then there is nothing to prove, so we can suppose that $X \neq \emptyset$.

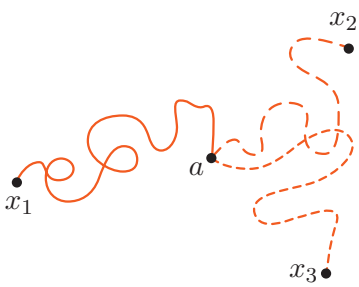

Figure 5.11

Fix some $a \in X$. If X is path-connected, then each point $x \in X$ can be joined to a by some path in X (Figure 5.11). Each such path defines a trail. The trail is the continuous image of the connected set $[0, 1]$ and so, by Lemma 3.6 the trail is a connected set, containing both a and x. The union of all these connected trails is the whole of X (as every point in X is the endpoint of at least one trail), and the intersection is non-empty since each trail contains the point a. It follows from Lemma 3.15 that the union is a connected set and so X is connected. ■

Corollary 5.14

The space $\big(C[0, 1], d_{\max}\big)$ is connected.

Proof This follows from Theorem 5.13 and Worked Exercise 5.12. ■

It follows from Theorem 5.13 that one way to show that a space is connected is to show that it is path-connected.

Worked Exercise 5.15

Let $A \subseteq \mathbb{R}^2$ be a countable set. Show that, for the induced Euclidean metric, $\mathbb{R}^2 - A$ is path-connected, and hence connected.

Solution

Since $\mathbb{R}^2$ contains a copy of $\mathbb{R}$, we know that $\mathbb{R}^2$ is uncountable. Also A is countable. Therefore we have $\mathbb{R}^2 - A$ uncountable (since $\mathbb{R}^2$ is the union of $\mathbb{R}^2 - A$ and A, and the union of two countable sets is countable by Book D, Chapter 16, Corollary 6.10).

Let $\mathbf{a}$ and $\mathbf{b}$ be distinct points of $\mathbb{R}^2 - A$. We describe how to find a path joining $\mathbf{a}$ and $\mathbf{b}$. Let L be the perpendicular bisector of the line segment joining $\mathbf{a}$ and $\mathbf{b}$. L is uncountable since it is a copy of $\mathbb{R}$. For each $\mathbf{x} \in L$, let $p_{\mathbf{x}} \colon [0, 1] \to \mathbb{R}^2$ be a path joining $\mathbf{a}$ to $\mathbf{b}$ with trail as illustrated in Figure 5.12. Let $\mathbf{x}$ and $\mathbf{y}$ be distinct points in L; then the trails of their corresponding paths $p_{\mathbf{x}}\big([0, 1]\big)$ and $p_{\mathbf{y}}\big([0, 1]\big)$ meet only at the points $\mathbf{a}$ and $\mathbf{b}$. In particular, $A \cap p_{\mathbf{x}}\big([0, 1]\big)$ is disjoint from $A \cap p_{\mathbf{y}}\big([0, 1]\big)$ whenever $\mathbf{x} \neq \mathbf{y}$. Since L is uncountable and A is countable, this means that there is $\mathbf{x} \in L$ for which $p_{\mathbf{x}}\big([0, 1]\big) \cap A = \emptyset$. For this $\mathbf{x}$, the map $p_{\mathbf{x}}$ is a path in $\mathbb{R}^2 - A$ joining $\mathbf{a}$ to $\mathbf{b}$. Thus $\mathbb{R}^2 - A$ is path-connected. It follows from Theorem 5.13 that $\mathbb{R}^2 - A$ is connected.

Countable sets were introduced in Book D, Chapter 16.

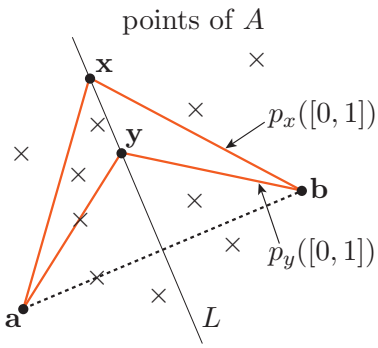

Figure 5.12

Remarks

1. This argument works for any $\mathbb{R}^n$ with $n > 1$ but fails for $\mathbb{R}$ (as we do not have enough room to construct L in one-dimensional space). In fact, as we know from Exercise 4.5, removing just one point from $\mathbb{R}$ gives a disconnected set!

2. Since the empty set is countable, this result confirms that $\mathbb{R}^2$ is path-connected.

Exercise 5.7

Show that $\mathbb{R}^2$ is not homeomorphic to $\mathbb{R}$.

Exercise 5.8

Let X be a set containing at least two points and let d_0 be the discrete metric on X. Show that X is not path-connected.

(*Hint*: use the result of Exercise 3.4.)

5.3 Products of path-connected spaces

We know that taking products preserves connectedness, so it is natural to ask whether taking products preserves path-connectedness. The answer is yes.

Theorem 5.16

Let (X, d_X) and (Y, d_Y) be non-empty metric spaces. The product space $(X \times Y, e)$, where e is a product metric, is path-connected if, and only if, (X, d_X) and (Y, d_Y) are both path-connected.

The proof is similar to that given for Theorem 4.8, and we omit it.

As with connectedness, this product result tell us that, for any intervals $I_1, \ldots, I_n$ of $\mathbb{R}$, $I_1 \times \cdots \times I_n$ is path-connected. In particular, if $I_k = \mathbb{R}$ for each k, we have another proof that $\mathbb{R}^n$ is path-connected.

We now turn our attention to unions and prove a result corresponding to Lemma 3.15 for connected spaces.

Theorem 5.17

Let (X, d) be a metric space, and let $\{A_i : i \in I\}$ be a family of path-connected subsets of X whose intersection is non-empty. Then $A = \bigcup_{i \in I} A_i$ is path-connected.

Proof Let $x, y \in A$; we want to find a path in A joining x and y.

Since $\bigcap_{i \in I} A_i \neq \emptyset$, we can fix a point $a \in \bigcap_{i \in I} A_i$ (Figure 5.13).

Since $x \in A$, there is an i so that $x \in A_i$. But $a \in A_i$ and so there is a path connecting x to a. Similarly, we can find a path joining y to a. We now use the result of Exercise 5.3 together with Lemma 5.4 to find a path joining x to y, proving that A is path-connected. ■

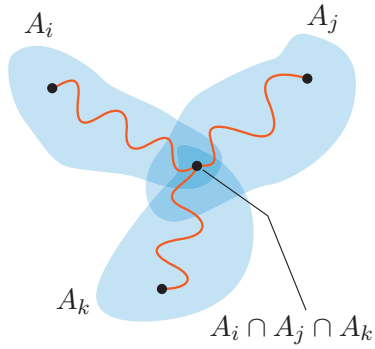

Figure 5.13

You may be wondering why we are so interested in path-connectedness when it seems so similar to connectedness. This is because there are differences between the two concepts. A path-connected space is always connected but *a connected space is not always path-connected*. In other words, just because a point inhabits the same connected space as another does not mean that it is possible to get from one to the other if you live in that space. In the next section we examine a space that is connected but not path-connected. In addition we will see that, unlike connectedness, path-connectedness is not preserved under closure.

6 The topologist's cosine

It seems reasonable to conjecture that connectedness implies path-connectedness, and that path-connectedness is preserved by closure. However, it is not possible to prove these conjectures because they are false! When it turns out to be difficult to *prove* a conjecture, then the next step is to look for a counter-example (and so *disprove* the conjecture). We do this in this section.

Constructing counter-examples is an important part of doing mathematics.

When constructing counter-examples it is often helpful to think about what could 'go wrong'. In this case, what are the possible obstructions to finding a path between two points in connected space? A path is a function from the closed interval $[0,1]$ to a space; the interval $[0,1]$ is not homeomorphic to $\mathbb{R}$, whereas $(0,1)$ is homeomorphic to $\mathbb{R}$. This suggests that, in some sense, $[0,1]$ cannot be 'stretched' as much as $(0,1)$. Intuitively, this is a consequence of the 'hard edges' of the closed interval: the image of $[0,1]$ under a continuous map needs to 'stop' somewhere, whereas the 'fuzzy' edges of $(0,1)$ can be smeared out indefinitely. We exploit this limitation of $[0,1]$ in our counter-example by constructing a connected space containing two points which are 'too far apart' to be reached via a path.

In the next chapter we investigate compactness and discover that $[0,1]$ is compact, whereas $(0,1)$ is not compact.

Our inspiration comes from the graph of the $\cos(\pi/x)$ function, which we investigated in Book D, Chapter 13, Worked Exercise 2.6. We defined this function $f\colon \mathbb{R} \to \mathbb{R}$ by

$$f(x) = \begin{cases} \cos(\pi/x), & \text{if } x \neq 0, \\ 0, & \text{if } x = 0. \end{cases}$$

In fact, it is not possible to redefine this function at 0 to make it continuous (with respect to the Euclidean metric).

Recall that this function, which oscillates very rapidly close to the origin, is continuous everywhere except at 0. The graph of f includes points arbitrarily close to each point on the line segment $\{0\} \times [-1,1]$ but only actually includes the point $(0,0)$ from this line segment.

If we consider the graph of the function f restricted to $(0,1]$, we have the set

$$E = \{(x, \cos(\pi/x))\colon 0 < x \leq 1\} \subset \mathbb{R}^2,$$

shown in Figure 6.1. This set (the graph E) is path-connected for its induced Euclidean metric. Once we specify any two points, a and b, on the

graph, we can bridge the gap between them using a continuous image of $[0, 1]$ (as there will always be finitely many oscillations between the specified points).

So we need to work a little harder to find our counter-example. We know that the graph does not include the points on the y-axis, so including these in our set may give us our counter-example. We also know that the graph gets arbitrarily close to the points on the y-axis, so it is likely that the closure of our graph, $\mathrm{Cl}_{(\mathbb{R}^2, d^{(2)})}(E)$, includes the vertical line segment $B = \{0\} \times [-1, 1]$ on the y-axis, shown in Figure 6.2. In fact, $\mathrm{Cl}_{(\mathbb{R}^2, d^{(2)})}(E) = E \cup B$.

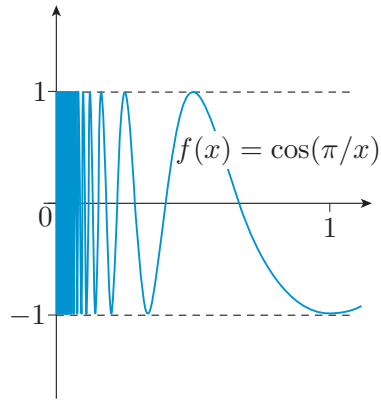

Figure 6.1

We know that connectedness is preserved under taking set closures so $\mathrm{Cl}_{(\mathbb{R}^2, d^{(2)})}(E)$ is connected. Is this new set path-connected? The answer is no. We obtain an idea as to why not if we try to visualise a trail between a point in B and a point in E. Our trail would need to cover 'infinitely many' large and dense oscillations, which is not possible.

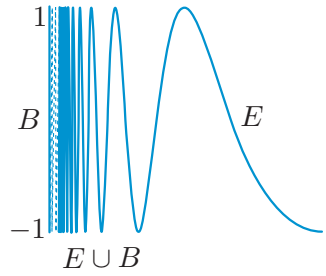

Figure 6.2

In order to simplify the calculations, our example, inspired by the $\cos(\pi/x)$ function, is constructed from straight line segments. This makes it much easier to determine whether a given point in the plane lies in the set.

Working out all the details in this example is a long process. Most of it is what pure mathematicians call *routine*; a word that should not be confused with *easy* or *elementary*. Routine work consists of providing details of standard pieces of analysis, such as proving the continuity of certain functions, or finding an upper bound for a set. We have concentrated on explaining the gist of the argument in the sketch proof that follows and have left some of the details as exercises. You may like to leave these exercises for now and return to them as part of your revision of Book F.

We begin by constructing a straight-line version of the set E we have just described.

For $n \in \mathbb{N}$, let A_n denote the set in the plane consisting of the line segment joining the point $(\frac{1}{n+1}, 0)$ to the point $(\frac{1}{2}(\frac{1}{n} + \frac{1}{n+1}), 1)$, together with the line segment joining $(\frac{1}{2}(\frac{1}{n} + \frac{1}{n+1}), 1)$ to $(\frac{1}{n}, 0)$, as shown in Figure 6.3.

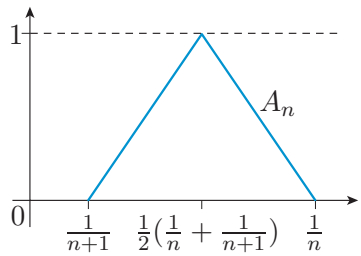

Figure 6.3

The algebraic description is unpleasant, but for future reference we give it:

$$A_n = \left\{(s, 2n(n+1)s - 2n) \colon \tfrac{1}{n+1} \leq s \leq \tfrac{1}{2}\left(\tfrac{1}{n} + \tfrac{1}{n+1}\right)\right\} \cup \left\{(s, 2(n+1) - 2n(n+1)s) \colon \tfrac{1}{2}\left(\tfrac{1}{n} + \tfrac{1}{n+1}\right) \leq s \leq \tfrac{1}{n}\right\}.$$

Let A be the plane subset

$$A = \bigcup_{n=1}^{\infty} A_n.$$

Then A is the *graph* of the function $l \colon (0, 1] \to \mathbb{R}$ given by

$$l(x) = \begin{cases} 2n(n+1)x - 2n, & \text{if } \frac{1}{n+1} < x \leq \frac{1}{2}\left(\frac{1}{n} + \frac{1}{n+1}\right), \ n \in \mathbb{N}, \\ 2(n+1) - 2n(n+1)x, & \text{if } \frac{1}{2}\left(\frac{1}{n} + \frac{1}{n+1}\right) < x \leq \frac{1}{n}, \ n \in \mathbb{N}, \end{cases}$$

This is similar to describing the familiar U-shaped quadratic curve as the graph of f, where $f(x) = x^2$, and the image set of F, where $F(x) = (x, x^2)$.

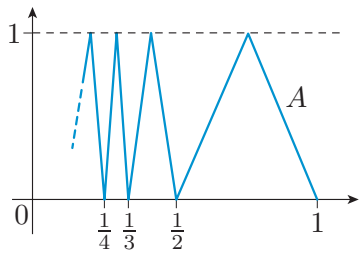

Figure 6.4

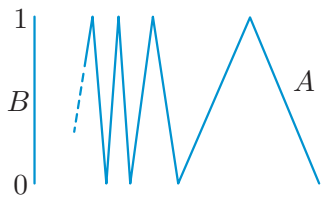

Figure 6.5

and shown in Figure 6.4. It is also the *image set* of the function

$$L\colon (0,1] \to \mathbb{R}^2 \text{ given by } L(x) = (x, l(x)).$$

We write d_A for $d_A^{(2)}$, the metric on A induced by the Euclidean metric on $\mathbb{R}^2$.

Let B be the vertical line segment joining the origin to the point $(0,1)$:

$$B = \{(0,y)\colon 0 \leq y \leq 1\}.$$

The *topologist's cosine* is the plane set

$$C = A \cup B,$$

consisting of all the points of A and B (Figure 6.5). We write d_C, for the metric on C induced by the Euclidean metric for the plane.

Lemma 6.1

The function $l\colon (0,1] \to \mathbb{R}$ is $(d_{(0,1]}^{(1)}, d^{(1)})$-continuous and the function $L\colon (0,1] \to \mathbb{R}^2$ is $(d_{(0,1]}^{(1)}, d^{(2)})$-continuous.

Proof We ask you to prove this in Exercise 6.1. ■

This allows us to show that A is a path-connected, and hence connected, subset of the plane, as follows.

Corollary 6.2

A is path-connected for d_A, the metric induced by the Euclidean metric on the plane.

Proof The set A is the image set of the path-connected set $(0,1]$ under the $(d_{(0,1]}^{(1)}, d^{(2)})$-continuous function L. So, by Theorem 5.8, A is path-connected. ■

We now know that the zigzag bit, A, of our set $C = A \cup B$ is d_A-connected. We also know that connectedness is preserved by closures (from Theorem 3.12). Therefore we can show that C is a d_C-connected subset of the plane by showing that $C = \mathrm{Cl}(A)$. To simplify what follows, we work in C and show that $C = \mathrm{Cl}_{(C,d_C)}(A)$ as this is sufficient to show that C is connected (by Theorem 3.12).

Theorem 6.3

C is d_C-connected, where d_C is the metric induced by the Euclidean metric on the plane.

Proof In order to show that C is d_C-connected, we show that the closure of A, when considered as a subset of C, is C.

Since the whole space in question is C, it is certainly the case that $\mathrm{Cl}_{(C,d_C)}(A) \subseteq C$. From Theorem 3.8 in Chapter 16 we know that $\mathrm{Cl}_{(C,d_C)}(A) = \{\mathbf{a} \in C : \mathbf{a} \text{ is a } d_C\text{-closure point of } A \text{ in } C\}$. So it remains only to show that each point $\mathbf{b}$ in $C - A = B$ is a d_C-closure point of A in C.

By Book D, Chapter 16, Worked Exercise 3.7, it is sufficient to find a sequence of points $(\mathbf{a}_n)$, with $\mathbf{a}_n \in A$ for each $n \in \mathbb{N}$ that is d_C-convergent to $\mathbf{b}$.

The point $\mathbf{b}$ is a point in B, and so can be written as $(0, b)$ for some b with $0 \leq b \leq 1$. We find a sequence of points, $(\mathbf{a}_n)$ in A, with limit $\mathbf{b}$, as demonstrated in Figure 6.6. In Exercise 6.2 we ask you to prove that $\mathbf{a}_n$ is d_C-convergent to $\mathbf{b}$ as $n \to \infty$. Hence $\mathbf{b} \in \mathrm{Cl}_{(C,d_C)}(A)$. Thus, since $\mathbf{b}$ was an arbitrary point in $C - A$, we have that $\mathrm{Cl}_{(C,d_C)}(A) = C$, and so C is d_C-connected by Theorem 3.12. ■

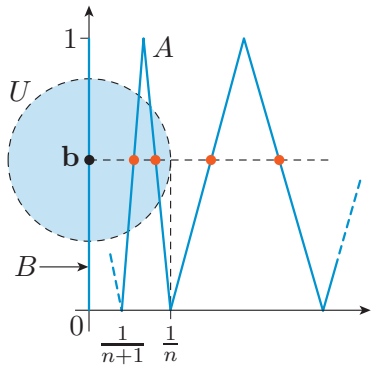

Figure 6.6

Thus we have a connected set, C, which is the closure of a path-connected set. It remains to show that C is not path-connected.

Theorem 6.4

C is not path-connected for the Euclidean metric.

Proof We have already observed (Corollary 6.2) that A is path-connected. Since B is also path-connected, the difficulty must lie in finding a path from a point in B to a point in A. The proof is by contradiction. All metrics are those induced by the relevant Euclidean metric on $\mathbb{R}^n$.

Suppose that C is path-connected; in particular, assume that there is a continuous function $P : [0, 1] \to C$ with initial point $P(0) = (0, 0)$ in B and final point $P(1) = (1, 0)$ in A, as shown in Figure 6.7. To describe the function P, we use the language of the path of a particle as a function of time. At time $t = 0$, the particle is at the point $(0, 0)$, at $t = 1$ it is at $(1, 0)$, and at time t it is at the point $P(t)$, somewhere in $\mathbb{R}^2$.

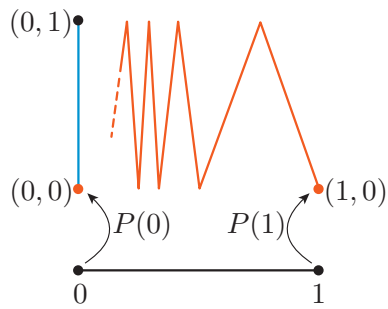

Figure 6.7

We used this language, with a snail instead of a particle, at the start of Section 4.

Write $P(t) = (P_1(t), P_2(t))$.

This introduces two new continuous functions, $P_1 : [0, 1] \to [0, 1]$ and $P_2 : [0, 1] \to [0, 1]$. In our particle–time language, $P_1(t)$ is the x-coordinate of the particle at time t, and $P_2(t)$ is its y-coordinate.

Now comes the key to the proof, certainly not an obvious idea. Since the particle starts at a point in the set B and ends up at a point in the set A, then there must be a *latest time* τ, *after which it never returns to* B. We find this time as follows.

If the particle kept popping back to B then it would not eventually come to rest in A at time $t = 1$.

Whenever the particle is in B, its x-coordinate is 0, so it is in B for those times t for which $P_1(t) = 0$, and for no other times. We collect all of these times together in the set $P_1^{-1}(\{0\}) \subseteq [0,1]$.

Note how this example depends crucially on the least upper bound property of $\mathbb{R}$.

The set $\{0\}$ is a $d^{(1)}$-closed subset of the x-axis and P_1 is continuous, so $P_1^{-1}(\{0\})$ is a $d^{(1)}$-closed subset of $[0,1]$. Since it is a subset of $[0,1]$ it is a bounded subset of $\mathbb{R}$, and so $P_1^{-1}(\{0\})$ has a least upper bound; call this number τ.

Note that $\tau < 1$, since at $t = 1$ the particle is known to be at the point $(1,0)$ in A and not in B.

Since $P_1^{-1}(\{0\})$ is $d^{(1)}$-closed, the supremum is contained in the set and so $P_1(\tau) = 0$. Therefore $P(\tau) \in B$ (as $\big(P_1(\tau), P_2(\tau)\big) = (0, P_2(\tau))$) and, for all later times t such that $1 \geq t > \tau$, we have $P(t) \in A$ (by our choice of τ).

To summarise: we have assumed that we have a path, $P(t) = \big(P_1(t), P_2(t)\big) \colon [0,1] \to C$, and have established that $\{P_1^{-1}\{0\}\}$ contains the supremum $\tau \in [0,1]$. We have deduced that $P_1(\tau) = 0$.

We now turn our attention to P_2. Remember that $P_1(t)$ is the x-coordinate of the path that we are trying to prove cannot exist because it oscillates so quickly. This oscillation is the key to our proof. We show that the y-value, $P_2(\tau)$, must simultaneously be every value in $[0,1]$ (which is, of course, impossible) and so gives a contradiction.

Since P_2 is continuous, for each $a \in [0,1]$, the set $P_2^{-1}(\{a\})$ is a $d^{(1)}$-closed subset of $[0,1]$. Therefore if, for some $a \in [0,1]$, we show that any $d^{(1)}$-open ball $B(\tau, r)$ intersects the $d^{(1)}$-closed set $P_2^{-1}(\{a\})$, then τ would be in the $d^{(1)}$-closure $\mathrm{Cl}(P_2^{-1}(\{a\}) = P_2^{-1}(\{a\})$; and so $P_2(\tau) = a$.

So, let $a \in [0,1]$. We claim that for each $r > 0$, there is $t \in B(\tau, r)$ for which $P_2(t) = a$.

To verify this claim, we fix $a \in [0,1]$, and let $B(\tau, r)$ be an open ball around τ. Then there is $\delta > 0$ such that

Note that we place no conditions on a or r (other than $r > 0$).

$$[\tau, \tau + \delta) \subset B(\tau, r).$$

We are looking for an element of $B(\tau, r)$ which P_2 maps to a so it suffices to find $s \in [\tau, \tau + \delta)$ for which $P_2(s) = a$.

We noted earlier that

$$0 = P_1(\tau) < P_1(\tau + \delta) \leq 1,$$

If $P_1(\tau + \delta) = 0$, then τ is not the least upper bound of $P_1^{-1}(\{0\})$.

and so we can find $N \in \mathbb{N}$ for which

$$0 = P_1(\tau) < \frac{1}{N} < P_1(\tau + \delta) \leq 1.$$

However, P_1 is continuous, and so, by the Intermediate Value Theorem (Theorem 4.6), there is $t_0 \in (\tau, \tau + \delta)$ for which

$$P_1(t_0) = \frac{1}{N},$$

as illustrated in Figure 6.8. As all of these points are in A, the way that we defined A ensures that $P_2(t_0) = 0$. Since

$$0 = P_1(\tau) < P_1(t_0) = \frac{1}{N} \quad \text{and} \quad \frac{1}{2}\left(\frac{1}{N} + \frac{1}{N+1}\right) < \frac{1}{N},$$

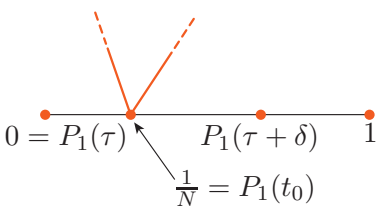

Figure 6.8

we can again use the Intermediate Value Theorem for P_1 to find $\tau < t_1 < t_0$ for which

$$P_1(t_1) = \frac{1}{2}\left(\frac{1}{N} + \frac{1}{N+1}\right),$$

and so, because of the way that A is defined, $P_2(t_1) = 1$ (see Figure 6.9).

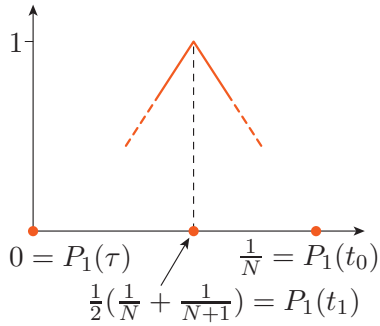

Figure 6.9

To summarise: for any $a \in [0,1]$ we want to show that any open ball around τ intersects $P_2^{-1}(\{a\})$. As $P_1(\tau) = 0$, and $P_1(\tau+\delta) > \frac{1}{N}$, we can use the Intermediate Value Theorem to find a t_0 with $P_1(t_0) = \frac{1}{N}$. We note that $P_2(t_0) = 0$. Here we are using the Intermediate Value Theorem to find a 'bottom corner' of the zigzag inside our r-ball around τ. We then use the same technique to find a 'top corner' of the zigzag inside our ε-ball, corresponding to the point t_1.

Thus $\tau < t_1 < t_0 < \tau + \delta$, $P_2(t_1) = 1$ and $P_2(t_0) = 0$. But P_2 is continuous, and so we can now use the Intermediate Value Theorem for P_2 to find $t_1 \leq s \leq t_0$ with $P_2(s) = a$. So we have proved that for each $B(\tau, r)$, there is $t \in B(\tau, r)$ for which $P_2(t) = a$ and so, since $P_2^{-1}(\{a\})$ is $d^{(1)}$-closed, we must have $\tau \in P_2^{-1}(\{a\})$. Therefore $P_2(\tau) = a$. However, *we put no conditions upon* a, other than to require $a \in [0,1]$. This gives our required contradiction, since we cannot have $P_2(\tau) = a$ for all $a \in [0,1]$!

To summarise: we deduce that $P_2(t_0) = 0$ and $P_2(t_1) = 1$ for two points, t_0 and t_1, inside our r-ball around τ. We then use the Intermediate Value Theorem again, this time applying it to P_2, to show that there must be a point s in $B(\tau, r)$ with $P_2(s) = a$. As every open ball around τ intersects the closed set $P_2^{-1}(\{a\})$, the point τ is a closure point of this set and so must be in the set, showing that $P_2(\tau) = a$ for any $a \in [0,1]$ which is impossible.

Hence C is not path-connected. ■

We have thus demonstrated that the topologist's cosine is a connected set that is not path-connected.

Exercise 6.1

Prove Lemma 6.1, which states that the function $l\colon (0,1] \to \mathbb{R}$ is $(d^{(1)}_{(0,1]}, d^{(1)})$-continuous and the function $L\colon (0,1] \to \mathbb{R}^2$ is $(d^{(1)}_{(0,1]}, d^{(2)})$-continuous.

Exercise 6.2

For $n \in \mathbb{N}$, let A_n denote the set in the plane consisting of the line segment joining the point $(\frac{1}{n+1}, 0)$ to the point $(\frac{1}{2}(\frac{1}{n} + \frac{1}{n+1}), 1)$, together with the line segment joining $(\frac{1}{2}(\frac{1}{n} + \frac{1}{n+1}), 1)$ to $(\frac{1}{n}, 0)$, as shown in Figure 6.3.

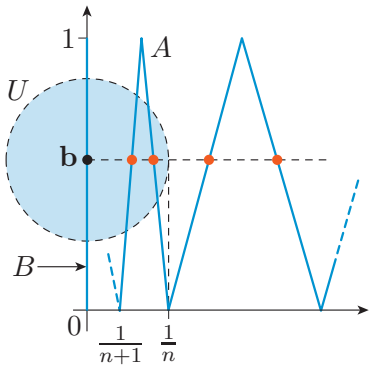

Figure 6.10

$$A_n = \left\{(s, 2n(n+1)s - 2n) \colon \tfrac{1}{n+1} \le s \le \tfrac{1}{2}\left(\tfrac{1}{n} + \tfrac{1}{n+1}\right)\right\}$$
$$\cup \left\{(s, 2(n+1) - 2n(n+1)s) \colon \tfrac{1}{2}\left(\tfrac{1}{n} + \tfrac{1}{n+1}\right) \le s \le \tfrac{1}{n}\right\}.$$

Let A be the plane subset

$$A = \bigcup_{n=1}^{\infty} A_n.$$

Let B be the vertical line segment joining the origin to $(0, 1)$:

$$B = \{(0, y) \colon 0 \le y \le 1\}.$$

Let $C = A \cup B$ with d_A, d_B and d_C being the metrics on A, b and C respectively induced by the Euclidean metric on $\mathbb{R}$. Let $\mathbf{b} = (0, b) \in B$, so $0 \le b \le 1$. Show that we can find a sequence of points in A with limit $\mathbf{b}$ as follows (see Figure 6.10). Let

$$x_n = \frac{b + 2n}{2n(n+1)}.$$

(a) Let l be as defined in Exercise 6.1. Show that $l(x_n) = b$ and hence that there is at least one point in A_n (namely $\mathbf{a}_n = (x_n, b)$) with y-coordinate equal to b.

(b) Let $r > 0$. Show that we can choose n so that the point $\mathbf{a}_n$ is inside $B_{d_C}(\mathbf{b}, r)$.

(c) Complete the proof that $\mathrm{Cl}_{(C,d_C)}(A) = C$.

Learning outcomes

We list those topics on which we may set assessment questions to test your understanding of the chapter.

After working through this chapter you should be able to:

(a) Define the terms *homeomorphism* and *topological invariant*.

(b) Understand the relationship between the closed and open sets in a metric space and a subspace of the metric space.

(c) Define the terms *disconnection*, *disconnected* and *connected*.

(d) Determine whether a given simple metric space is connected.

(e) Produce disconnections of some disconnected spaces.

(f) Specify which subsets of $\mathbb{R}$ are connected.

(g) Appreciate the generalisation of the Intermediate Value Theorem.

(h) Explain what are meant by a *path* and a *path-connected* space.

(i) Show that certain spaces are path-connected.

(j) State and make use of properties of path-connected spaces.

(k) Appreciate the links and differences between path-connected and connected spaces.

(l) Describe an example of a metric space that is connected, but not path-connected.

Solutions to exercises

Solution to Exercise 1.1

One homeomorphism from A to B is the function $f\colon A \to B$ given by $f\colon x \mapsto nx$ (which is clearly a $(d_A^{(1)}, d_B^{(1)})$-continuous bijection). The number n does not have to be a natural number, but n must not be 0 since we cannot have a bijection between a point and an interval.

Solution to Exercise 1.2

One such map is the linear map that takes a to c and b to d and stretches (or shrinks) the first interval to cover the second. This is the function $f_I\colon X \to Y$ given by $f_I\colon x \mapsto c + \frac{d-c}{b-a}(x-a)$. As f_I is a linear map, it is clearly a bijective $(d_X^{(1)}, d_Y^{(1)})$-continuous map with a $(d_Y^{(1)}, d_X^{(1)})$-continuous inverse.

It does not matter if $a < b$ and $c > d$ since the homeomorphism can 'flip the interval over'. This action is still continuous. For example, $f\colon x \mapsto -x$ gives a homeomorphism from $[0,1]$ to $[-1,0]$.

Solution to Exercise 1.3

The map $f_I\colon x \mapsto c + \frac{d-c}{b-a}(x-a)$ used in Exercise 1.2, when restricted to the half-open interval, gives a homeomorphism from X to Y. The intervals $(a,b]$ and $[c,d)$ are homeomorphic since the map $f\colon x \mapsto -x$ gives a homeomorphism from $[a,b)$ to $(-b,-a]$. So any half-open interval of the form $[a,b)$ can be 'flipped over' to one of the form $(x,y]$ before applying a homeomorphism of the form f_I as described earlier.

Solution to Exercise 1.4

Solution to Exercise 1.5

The interval $(a,b]$ incudes b. Therefore we must have $f(b) \in (c,d)$. However, $f(b) = d$ and $d \notin (c,d)$ so f is not a map from $(a,b]$ to (c,d).

Solution to Exercise 1.6

Since one set is infinite and the other is finite, there can be no bijection between them and so no homeomorphism.

Solution to Exercise 1.7

Consider the function $\phi\colon \mathbb{R} \to (a,\infty)$ given by $\phi(x) = a + \exp(x)$. This is a sum of basic continuous functions (so is continuous) and has a continuous inverse (namely $\phi^{-1}\colon (a,\infty) \to \mathbb{R}$ given by $\phi^{-1}(x) = \log(x-a)$). The fact that the inverse is defined shows that ϕ is onto, and standard properties of $\exp(x)$ show that it is one–one.

Solution to Exercise 2.1

There are many possible solutions. One such is $X = (0,1) \cup \{2\}$, with the metric induced by the Euclidean metric on $\mathbb{R}$, $d_X^{(1)}$. For then $(0,1) = [0,1] \cap \big((0,1) \cup \{2\}\big)$ and so, by Theorem 3.12 of Chapter 16, $(0,1)$ is $d_X^{(1)}$-closed in $X = (0,1) \cup \{2\}$.

Solution to Exercise 2.2

Note that these solutions are not unique.

(a) We are working in the subspace (X,d) where $X = (-\infty,0) \cup (1,\infty)$ and d is the metric induced by the Euclidean metric on $\mathbb{R}$. Therefore we know that the d-closed sets in X will be the intersections of $d^{(1)}$-closed sets in $\mathbb{R}$ with X. So to show that $(-\infty,0)$ is d-closed, we need to find a $d^{(1)}$-closed set A in $\mathbb{R}$ such that $A \cap X = (-\infty,0)$. Now $(-\infty,\frac{1}{2}]$ is $d^{(1)}$-closed in $\mathbb{R}$ and $(-\infty,0) = (-\infty,\frac{1}{2}] \cap X$, and so is d-closed when considered as a subset of X. (We could have used any set $(-\infty,x]$ with $x \in [0,1)$.) Also $(-\infty,0) = (-\infty,\frac{1}{2}) \cap X$ and so is d-open (as $(-\infty,\frac{1}{2})$ is $d^{(1)}$-open in $\mathbb{R}$).

(b) Here we are working in the subspace (X,d) where $X = (-\infty,0] \cup (1,\infty)$ and d is the induced metric by the Euclidean metric on $\mathbb{R}$. Therefore we know that the d-closed sets in X will be the intersections of $d^{(1)}$-closed sets in $\mathbb{R}$ with X. The interval $(-\infty,0] = (-\infty,\frac{1}{2}] \cap X$ and so is d-closed (as $(-\infty,\frac{1}{2}]$ is $d^{(1)}$-closed in $\mathbb{R}$). Similarly, $\{x \in \mathbb{R} \colon -\infty < x \leq 0\} = (-\infty,\frac{1}{2}) \cap X$ and so is d-open (as $(-\infty,\frac{1}{2})$ is $d^{(1)}$-open in $\mathbb{R}$).

(c) This time we are working in (X,d) where $X = \{0\} \cup (3,4]$ and d is the metric induced by the Euclidean metric on $\mathbb{R}$, and so we need to consider the intersections of $d^{(1)}$-closed and $d^{(1)}$-open subsets of $\mathbb{R}$ with X. The interval $[-1,1]$ is $d^{(1)}$-closed in $\mathbb{R}$ and the intersection of this interval with X is $\{0\}$, and so $\{0\}$ is d-closed in X. Also $\{0\} = (-1,1) \cap X$ and so is d-open in X (as $(-1,1)$ is $d^{(1)}$-open in $\mathbb{R}$).

(d) First we want to show that $(0,1)$ is d-closed as a subset of $\mathbb{R} - \mathbb{Z}$ where d is the metric induced by the Euclidean metric on $\mathbb{R}$. As we are working with the induced metric, we need to consider the intersections of $d^{(1)}$-closed sets in $\mathbb{R}$ with $\mathbb{R} - \mathbb{Z}$. Now $[0,1]$ is $d^{(1)}$-closed in $\mathbb{R}$ and so $(0,1) = [0,1] \cap (\mathbb{R} - \mathbb{Z})$ is d-closed in $\mathbb{R} - \mathbb{Z}$. Also $(0,1) = (0,1) \cap (\mathbb{R} - \mathbb{Z})$ and so is d-open in $\mathbb{R} - \mathbb{Z}$.

Solution to Exercise 3.1

(a) Here $X = (-\infty, 0) \cup (1, \infty)$ and d_X is the metric induced by the Euclidean metric on $\mathbb{R}$. We need to find two non-empty disjoint sets, each of which is simultaneously d_X-closed and d_X-open as a subset of (X, d_X), and whose union is the whole of X. We make the natural choices, $U = (-\infty, 0)$ and $V = (1, \infty)$, which we know are each simultaneously d_X-closed and d_X-open from Exercise 2.2. Therefore a disconnection is $\{(-\infty, 0), (1, \infty)\}$.

(b) Here $X = (-\infty, 0] \cup (1, \infty)$ and d_X is the metric induced by the Euclidean metric on $\mathbb{R}$. We need to find two non-empty disjoint sets, each of which is simultaneously d_X-closed and d_X-open as a subset of (X, d_X), and whose union is the whole of $(-\infty, 0] \cup (1, \infty)$. We make the natural choices, $U = (-\infty, 0]$ and $V = (1, \infty)$, which we know are each simultaneously d_X-closed and d_X-open from Exercise 2.2. Therefore a disconnection is $\{(-\infty, 0], (1, \infty)\}$.

(c) Here $X = \{0\} \cup (3, 4]$ and d_X is the metric induced by the Euclidean metric on $\mathbb{R}$. We need to find two non-empty disjoint sets, each of which is simultaneously d_X-closed and d_X-open as a subset of (X, d_X), and whose union is the whole of $\{0\} \cup (3, 4]$. We make the natural choices, $U = \{0\}$ and $V = (3, 4]$, which we know are each simultaneously d_X-closed and d_X-open from Exercise 2.2. Therefore a disconnection is $\{\{0\}, (3, 4]\}$.

Solution to Exercise 3.2

Since the only subset of $\emptyset$ is $\emptyset$ itself, it follows that there is only one possible subset – namely, $\emptyset$ – and so no disconnection exists. So this space is connected.

Solution to Exercise 3.3

The only subsets of X are $\emptyset$, $\{x\}$ and X. Thus we cannot find two disjoint proper subsets and no disconnection of X can exist. Therefore X is connected.

Solution to Exercise 3.4

We look for a proper subset which is simultaneously d_0-closed and d_0-open. Since d_0 is the discrete metric, we know from Chapter 16 that all subsets of X are d_0-closed. In particular, each of $\{x\}$ and $\{x\}^c$ is d_0-closed. As $\{x\}^c$ is d_0-closed $\{x\} = (\{x\}^c)^c$, being the complement of a d_0-closed set, is d_0-open. As X contains at least two points, we know that $\{x\} \neq X$. Thus, the set $\{x\}^c$ is non-empty. Therefore $\{x\}$ is a d_0-closed and d_0-open proper subset of X, and so X is disconnected by Theorem 3.2.

Solution to Exercise 3.5

If X is disconnected then, by definition, X can be expressed as the union of two non-empty d-closed sets E and F. Therefore we need to show that if X can be expressed as the disjoint union of two d-closed sets $X = E \cup F$ with $E \cap F = \emptyset$, then X can be expressed as the disjoint union of two non-empty sets, each of which is simultaneously d-closed and d-open. Assume that $X = E \cup F$ with E and F disjoint, non-empty and d-closed. Then $E = X - F$ and so E is the complement of a d-closed set and is d-open. Similarly, F is d-open as it is the complement of E. Therefore E and F are each simultaneously d-closed and d-open and X is disconnected.

Solution to Exercise 3.6

We look for a way to divide $\mathbb{R} - \mathbb{Q}$ into two disjoint, $d_{\mathbb{R}-\mathbb{Q}}$-closed subsets of $\mathbb{R} - \mathbb{Q}$. As $\mathbb{R} - \mathbb{Q}$ has the metric induced by the Euclidean metric on $\mathbb{R}$, we know that a $d_{\mathbb{R}-\mathbb{Q}}$-closed subset of $\mathbb{R} - \mathbb{Q}$ is the intersection of a $d^{(1)}$-closed subset of $\mathbb{R}$ with $\mathbb{R} - \mathbb{Q}$.

Let q be a rational number, such as 42. Then $E = (-\infty, 42] \cap (\mathbb{R} - \mathbb{Q})$ and $F = [42, \infty) \cap (\mathbb{R} - \mathbb{Q})$ are two disjoint $d_{\mathbb{R}-\mathbb{Q}}$-closed sets whose union is $\mathbb{R} - \mathbb{Q}$. Therefore by Proposition 3.4 $(\mathbb{R} - \mathbb{Q}, d_{\mathbb{R}-\mathbb{Q}})$ is disconnected.

Solution to Exercise 3.7

There are many possible solutions to this question. One is as follows. Let $X = \{0, 1\}$ with the discrete metric, which is disconnected by Exercise 3.4. Let $Y = \{1\}$ with the discrete metric, which is connected by Exercise 3.3. Define $f\colon X \to Y$ by $f(x) = 1$ for $x \in X$. Then f is (d, e)-continuous (by Proposition 1.10). Also any discrete space with more than one point is disconnected (by Exercise 3.4) and so f is an onto (d_0, d_0)-continuous map from the disconnected space (X, d_0) to the connected space (Y, d_0). (In fact, any map from a discrete space with at least two points onto a connected space will suffice here.)

Solution to Exercise 3.8

Suppose that χ_U is (d, d_0)-continuous. The set $\{1\}$ is both d_0-closed and d_0-open in $\{0, 1\}$. Therefore $\chi_U^{-1}(\{1\}) = U$ is d-closed and d-open in X (by Book D, Chapter 16, Theorem 7.1 and Corollary 7.2).

Solution to Exercise 3.9

Suppose that U is both d-closed and d-open. We need to show that the preimage of any d_0-closed subset of $\big(\{0, 1\}\big), d_0)$ is d-closed. Since d_0 is the discrete metric, every subset of the set is d_0-closed; and since it is a finite set, we can simply list the possible subsets and deal with each in turn. The subsets are $\emptyset, \{0\}, \{1\}, \{0, 1\}$. Now $\chi_U^{-1}(\emptyset) = \emptyset$, which is d-closed in X; $\chi_U^{-1}(\{0\}) = U^c$, which is d-closed in X; $\chi_U^{-1}(\{1\}) = U$ and $\chi_U^{-1}(\{0, 1\}) = X$, which are both d-closed sets in X. Therefore χ_U is a (d, d_0)-continuous map by Book D, Chapter 16, Theorem 7.1.

Solution to Exercise 3.10

There are (infinitely) many ways to answer this question. An example is: $\mathbb{Q} = \{q \in \mathbb{Q} : q < \sqrt{2}\} \cup \{q \in \mathbb{Q} : q > \sqrt{2}\}$ and $\mathbb{Q} = \{q \in \mathbb{Q} : q < \pi\} \cup \{q \in \mathbb{Q} : q > \pi\}$.

Solution to Exercise 3.11

We write d for the induced metric $d^{(1)}_{(\mathbb{R}-\mathbb{Z})}$. By Theorem 2.1 the interval $I_n = (n, n+1)$, where $n \in \mathbb{Z}$, is d-open in X (as it is $d^{(1)}$-open in $\mathbb{R}$ and equal to the intersection of itself with $\mathbb{R}$). It is also d-closed (as it is equal to the intersection of the $d^{(1)}$-closed interval $[n, n+1]$ with X). Therefore $I_n = (n, n+1)$ is simultaneously d-open and d-closed, so cannot be a proper subset of a connected subset of X (by Theorem 3.2). By the given assumption, I_n is connected in X. Hence the intervals I_n form the components of X.

Solution to Exercise 3.12

When $(0, 1)$ has the discrete metric, each of the sets $\{x\}$, where $x \in (0, 1)$, is simultaneously d_0-open and d_0-closed (as every subset of any space with the discrete metric is simultaneously d_0-closed and d_0-open). Therefore any subset of $(0, 1)$ containing more than one element is disconnected (since if x and y are both in the same subset, A, then $\{x\}$ and $A - \{x\}$ would form a disconnection) and so the components must be the singleton sets $\{x\}$.

Solution to Exercise 3.13

Let (X, d) be a metric space, and let $x \in X$. Write $\mathrm{Cl}(X)$ for $\mathrm{Cl}_{(X,d)}(X)$.

It follows from Theorem 3.12 that $\mathrm{Cl}(C_d(x))$ is connected. Now $\mathrm{Cl}(C_d(x))$ contains $C_d(x)$. Also since $\mathrm{Cl}(C_d(x))$ is connected we have $\mathrm{Cl}(C_d(x)) \in \mathcal{F}_x = \{A \subseteq X : x \in A \text{ and } A \text{ is connected}\}$. Therefore $\mathrm{Cl}(C_d(x)) \subseteq \bigcup_{A \in \mathcal{F}_x} A$, which is equal to $C_d(x)$. So $\mathrm{Cl}(C_d(x)) = C_d(x)$, and hence $C_d(x)$ is closed.

Solution to Exercise 4.1

Let $B = \{-a : a \in A\}$. Since A is bounded below, B is bounded above, and so B has a least upper bound, M say. Since $b \leq M$ for all $b \in B$, it follows that $-a \leq M$ for all $a \in A$, and so $-M \leq a$ for all $a \in A$. Hence $-M$ is a lower bound for A.

If $m' > -M$, then $-m' < M$, and so there is $b \in B$ for which $-m' < b$. But this means there is $a \in A$ with $-m' < -a$, and this rearranges to give $a < m'$. Hence $-M$ is the greatest lower bound of A.

Solution to Exercise 4.2

By Worked Exercise 1.13, the interval $(0, 1)$ is homeomorphic to $\mathbb{R}$. By Corollary 3.7, connectedness is preserved by homeomorphism. Therefore $(0, 1)$ must be connected.

Solution to Exercise 4.3

The interval $(0, 1)$ with the Euclidean metric is connected. Therefore there is a single component, namely $(0, 1)$ itself.

Solution to Exercise 4.4

All these intervals contain (a, b) and are contained in the interval $[a, b]$, the $d^{(1)}$-closure of (a, b) in $\mathbb{R}$. Therefore they are connected, by Theorem 3.12.

Solution to Exercise 4.5

A disconnection of $\mathbb{R} - \{a\}$ is $\{(-\infty, a), (a, \infty)\}$.

Solution to Exercise 4.6

The continuous image of a connected set is connected by Lemma 3.6. The map $f : x \mapsto \big(\sin(x), \cos(x)\big)$ gives a $(d_I^{(1)}, d_C^{(2)})$-continuous map from the interval $I = [0, 2\pi)$ to C. Therefore C is connected.

Solution to Exercise 4.7

The proof is by contradiction. We know that A and B are both connected. Suppose that there is a homeomorphism $\phi\colon A \to B$. Let $a \in A$ be the point for which $\phi(a) = 0$ and $b \in A$ be the point for which $\phi(b) = 1$. The restriction rule for homeomorphisms implies that $A - \{a, b\}$ is homeomorphic to $\phi\big(A - \{a, b\}\big) = B - \{0, 1\} = (0, 1)$. This, however, is a contradiction since $A - \{a, b\}$ is disconnected, while $(0, 1)$ is connected. Thus A is not homeomorphic to B.

Note that we need to consider both the point a for which $\phi(a) = 0$ and the point b for which $\phi(b) = 1$ to ensure a disconnection here. This is because it is possible that $\phi(1) = 0$ or $\phi(1) = 1$, in which case $A - \{1\} = (0, 1)$, which is connected. The second point ensures that we have a point interior to $(0, 1)$ with which to disconnect A.

Solution to Exercise 4.8

Let (X, d_X) and (Y, d_Y) be non-empty metric spaces, and suppose that $X \times Y$ is e-connected. We know that the projection functions are continuous, and that $X = p_X(X \times Y)$, $Y = p_Y(X \times Y)$. So X and Y are images of a connected space under continuous mappings, and so are connected (by Lemma 3.6).

Solution to Exercise 4.9

The function f is continuous and the space $(\mathbb{R}^2, d^{(2)})$ is connected, so we can apply Theorem 4.6, the Intermediate Value Theorem. We must find a point in $\mathbb{R}^2$ at which the function takes a value less than 42 and another point at which the function takes a value greater than 42. For example,

$$f(0, 0) = 0 \quad \text{and} \quad f(2, 0) = 5 \times 2^5 = 160.$$

So, by the Intermediate Value Theorem, there is a point in the plane at which f attains the value 42. (The theorem does not tell us what this point is, just that such a point exists, since the Intermediate Value Theorem is an existence result.)

See the discussion after Book D, Chapter 13, Worked Exercise 2.13.

Solution to Exercise 5.1

By Theorem 7.1 of Chapter 16, the constant function $p\colon [0,1] \to X$ given by $p(t) = x$ will be $(d^{(1)}_{[0,1]}, d)$-continuous if, and only if, $p^{-1}(E)$ is $d^{(1)}_{[0,1]}$-closed whenever E is a d-closed set in X. So let E be a d-closed set in X. Then either $x \in E$ or $x \notin E$. If $x \in E$ then $p^{-1}(E) = [0,1]$, which is $d^{(1)}_{[0,1]}$-closed. If $x \notin E$ then $p^{-1}(E) = \emptyset$, which is also $d^{(1)}_{[0,1]}$-closed. Therefore we have the result.

Solution to Exercise 5.2

Define $f\colon [0,1] \to X$ by $f(t) = (t, \frac{1}{2}t)$. We note that $(p_1 \circ f)(t) = t$ and $(p_2 \circ f)(t) = \frac{1}{2}t$ are $(d^{(1)}_{[0,1]}, d^{(1)}_{[0,1]})$-continuous functions from $[0,1]$ to $[0,1]$. So, by Theorem 3.12 of Chapter 15, the map f is $(d^{(1)}_{[0,1]}, d^{(2)}_X)$-continuous. Hence, since $f(0) = (0,0)$ and $f(1) = (1, \frac{1}{2})$, the map f gives a path joining $(0,0)$ and $(1, \frac{1}{2})$.

Solution to Exercise 5.3

In this solution, all metrics are the metrics induced by the Euclidean metric on $\mathbb{R}$.

Define $q\colon [0,1] \to X$ by $q(t) = p(1-t)$. Then $q(0) = p(1) = b$ and $q(1) = p(0) = a$. Moreover, since $q = p \circ f$, where $f\colon [0,1] \to [0,1]$ is the basic continuous function given by $f(t) = 1 - t$, it follows that q is a composition of continuous functions and so is continuous. Hence q is a path in X with initial point b and final point a.

Solution to Exercise 5.4

If $c \neq d$ are two points in $[a,b]$ then we need to show that there is a path within $[a,b]$ from c to d.

If $a = b$, we are done. Assume (without loss of generality) that $a < b$ and let $c, d \in [a,b]$. If $c = d$, we are done, so we assume (without loss of generality) that $c < d$. We define $P\colon [0,1] \to [a,b]$ by

$$P(t) = (1-t)c + td.$$

Note that
$(1-t)c + td \leq (1-t)b + tb = b$
and
$(1-t)c + td \geq (1-t)a + ta = a$.

We know from Worked Exercise 5.6 that P defines a path from c to d in $\mathbb{R}$. Thus, using Theorem 1.6 in Chapter 15, we can deduce that P defines a path in $[a,b]$ if $P([0,1]) \subseteq [a,b]$. This is clearly true 'by inspection'. Therefore $[a,b]$ is path-connected.

Solution to Exercise 5.5

The metrics have been added in to the relevant section below.

Continuity: let $p_j\colon \mathbb{R}^n \to \mathbb{R}$ be the jth projection function, defined by $p_j(\mathbf{a}) = a_j$. Then $(p_j \circ P)(t) = (1-t)a_j + tb_j$, so $p_j \circ P$ is linear and so is clearly $(d^{(1)}_{[0,1]}, d^{(1)})$-continuous. Since $p_j \circ P$ is $(d^{(1)}_{[0,1]}, d^{(1)})$-continuous for $j = 1, 2, \ldots, n$, it follows from the generalisation of Theorem 3.12 of Chapter 15 that P is $(d^{(1)}_{[0,1]}, d^{(n)})$-continuous.

Solution to Exercise 5.6

For any $x \in [0,1]$, and any $s, t \in [0,1]$,

$$\begin{aligned} |P(t)(x) - P(s)(x)| &= |(1-t)f(x) + tg(x) - (1-s)f(x) - sg(x)| \\ &= |(t-s)(g(x) - f(x))| \\ &= |t-s|\,|g(x) - f(x)| \\ &\leq |t-s| d_{\max}(f,g). \end{aligned}$$

Taking the maximum over $x \in [0,1]$, we have

$$d_{\max}(P(t), P(s)) \leq |t-s| d_{\max}(f,g).$$

Therefore P is a $(d^{(1)}_{[0,1]}, d_{\max})$ Lipschitz function and so is $(d^{(1)}_{[0,1]}, d_{\max})$-continuous (by Theorem 3.2 of Chapter 15).

Solution to Exercise 5.7

Suppose that there is a homeomorphism $\phi\colon \mathbb{R} \to \mathbb{R}^2$. Then removing a point, p, from $\mathbb{R}$ and removing $\phi(p)$ from $\mathbb{R}^2$ would give us a homeomorphism from $\mathbb{R} - \{p\}$ to $\mathbb{R}^2 - \{\phi(p)\}$. This is impossible since $\mathbb{R} - \{p\}$ is not connected (and so is not path-connected), but $\mathbb{R}^2 - \{p\}$ is path-connected (by Worked Exercise 5.15) and path-connectedness is a topological invariant by Corollary 5.9.

Solution to Exercise 5.8

We know, from Exercise 3.4, that X is not connected. It follows from Theorem 5.13 that X cannot be path-connected.

Solution to Exercise 6.1

For each n, let $l_{n,1}$ be the map l restricted to the set: $I_{n,1} = \left\{x \in \mathbb{R}\colon \frac{1}{n+1} \leq x \leq \frac{1}{2}\left(\frac{1}{n} + \frac{1}{n+1}\right)\right\}$ and $l_{n,2}$ be the map l restricted to the set $I_{n,1} = \left\{x \in \mathbb{R}\colon \frac{1}{2}\left(\frac{1}{n} + \frac{1}{n+1}\right) \leq x \leq \frac{1}{n},\ n \in \mathbb{N}\right\}$. Both maps $l_{n,i}$ are linear and so are $(d^{(1)}_{I_{n,i}}, d^{(1)})$-continuous. Also the maps $l_{n,i}$ agree on appropriate points. Therefore l is continuous by Lemma 5.3.

Let p_1 and p_2 be the projection functions from $\mathbb{R}^2$ to $\mathbb{R}$. Then $(p_1 \circ L)(x) = x$ and $(p_2 \circ L)(x) = l(x)$, and so $p_1 \circ L$ and $p_2 \circ L$ are both $(d^{(1)}_{[0,1]}, d^{(1)})$-continuous functions. Hence L is $(d^{(1)}_{[0,1]}, d^{(2)})$-continuous, by Proposition 1.19 of Chapter 14.

Solution to Exercise 6.2

(a) Since $0 \leq b \leq 1$, we have:

$$\frac{1}{n+1} \leq x_n \leq \frac{2n+1}{2n(n+1)} = \frac{1}{2}\left(\frac{1}{n} + \frac{1}{n+1}\right),$$

and so $l(x_n) = 2n(n+1)x_n - 2n = (b+2n) - 2n = b$.

(b) We want $\mathbf{a}_n$ to be inside $B_{d_C}(\mathbf{b}, r)$. This occurs if the x-coordinate, x_n is less than r.

$$x_n = \frac{b+2n}{2n(n+1)}$$

so we need to find N with

$$\frac{b+2N}{2N(N+1)} < r.$$

Choosing any N such that $2N((N+1)r-1)) > b$, which always exists, will suffice.

(c) Part (a) shows that we can find a sequence of points $(\mathbf{a}_n)$, where $\mathbf{a}_n \in A_n$ and where the x-coordinate of each point is b. Part (b) shows that for any $r > 0$ we can find an $N \in \mathbb{N}$ such that $\mathbf{a}_n \in B_{d_C}(\mathbf{b}, r)$. In other words, we have shown that $\mathbf{a}_n$ is d_C-convergent to $\mathbf{b}$. Therefore $\mathbf{b}$ is a limit point of $(\mathbf{a}_n)$ and $\mathbf{b} \in \mathrm{Cl}_{(C,d_C)}(A)$. Thus $\mathrm{Cl}_{(C,d_C)}(A) = C$.

Chapter 22

Compactness

Introduction

We are interested in studying continuous functions. In Book D, Chapter 13, just after Example 2.17, we noted that the Intermediate Value, Boundedness and Extreme Value Theorems each hold for all continuous functions on all $d^{(1)}$-closed and $d^{(1)}$-bounded intervals in $\mathbb{R}$. This suggests that they are the result of some kind of interaction between continuity and closed and bounded intervals. In Chapter 21, we saw that the Intermediate Value Theorem can be generalised to apply to any connected metric space. In this chapter, we generalise the idea of 'closed and bounded'-ness to metric spaces, enabling us to generalise the Extreme Value Theorem.

Recall that connectedness, the characteristic of a space which allows the Intermediate Value Theorem to hold, gives an intuitive idea as to whether or not a metric space is in one 'piece'. The characteristic which enables the Extreme Value Theorem to hold turns out to be *(sequential) compactness*, which gives us another idea of 'size' or 'stretchability' for a space.

We stated, but did not prove, the Boundedness Theorem and the Extreme Value Theorem in Chapter 13.

Boundedness Theorem

Let $f\colon [a,b] \to \mathbb{R}$ be a continuous function on $[a,b]$. Then f is bounded on $[a,b]$. That is, there is a real number M such that $|f(x)| \leq M$ for each $x \in [a,b]$.

Here 'continuous' means $(d^{(1)}_{[a,b]}, d^{(1)})$-continuous.

Extreme Value Theorem

Let $f\colon [a,b] \to \mathbb{R}$ be a continuous function on $[a,b]$. Then there are numbers $c, d \in [a,b]$ such that

$$f(c) \leq f(x) \leq f(d) \quad \text{for each } x \in [a,b].$$

Informally, the Extreme Value Theorem tells us that a $(d^{(1)}_{[a,b]}, d^{(1)})$-continuous function on a $d^{(1)}$-closed and bounded interval of the real line is bounded and attains its bounds. Formally, this theorem tells us that two numbers, n and m, exist, with $n \leq f(x) \leq m$ for each $x \in [a,b]$, (the function is bounded) and that two numbers c and d exist, in $[a,b]$, with $f(c) = n$ and $f(d) = m$ (the function attains it bounds).

Note that we can have $n = m$ if f is a constant function.

In the first section of this chapter, we prove this result for the interval $[a,b]$, and pull out from the proof the key characteristic of $d^{(1)}$-closed and bounded intervals that is necessary for the proof to 'work'. This turns out to be the rather remarkable fact that *every* sequence in $[a,b]$ has a convergent subsequence, and it is this characteristic that gives us the definition of sequential compactness.

In the second section we formally define sequential compactness, and prove the Heine–Borel Theorem, an important result in real analysis which characterises all sequentially compact subsets of $\mathbb{R}^n$. We then, in Section 3, turn our attention to the function space $C[0,1]$ and look at the sequentially compact subsets of this space.

In Section 4 we look at compactness, a notion of size that is defined using the open set structure of a space. We go on to prove that the definitions of sequential compactness and of compactness are equivalent for all metric spaces.

Chapter guide

The ideas introduced within this chapter are subtle, and you may find that it takes some effort and practice before you are fully comfortable with them. This is normal – the material is often regarded as quite difficult.

The most important ideas of this unit appear in Sections 2 and 4. You will probably need to spend most of your time studying these two sections.

In Section 2 we introduce the idea of sequential compactness and investigate how sequences and subsequences can be used to encapsulate various properties of continuous functions on closed and bounded intervals. Later, in Section 4, we see how the ideas of an *open cover* and a *finite subcover* can be used to give an alternative approach. We show that these two approaches both lead to the same destination for a metric space in Section 5. The proof that the two notions of compactness are equivalent for metric spaces is long. You should focus on following the general ideas behind the proofs and not spend too much time studying the technical details. Section 3 looks at sequentially compact subsets of the function space $C[0,1]$ and gives a relatively simple technique that we can use to verify whether a given subset of $C[0,1]$ is sequentially compact. If you are short of time when reading this section, then you should focus on ensuring that you can use this technique.

Notation

In this chapter, when considering a subset X, of $\mathbb{R}^n$, with the metric, $d_X^{(n)}$, induced by the usual Euclidean metric on $\mathbb{R}^n$, we will often drop the subscript and use the notation $d^{(n)}$ to refer to $d_X^{(n)}$.

1 Two important theorems

In this section we prove the Boundedness Theorem and the Extreme Value Theorem for subsets of the real line (with the Euclidean metric). The Boundedness Theorem is a consequence of the Extreme Value Theorem, so we concentrate our efforts on finding a generalisation of the latter result. The Extreme Value Theorem applies to closed and bounded intervals of the real line, and we examine the proof of this result to try to pinpoint the precise aspect of 'closed and bounded'-ness upon which the proof relies.

Remark

Note that in the following discussion we use the word 'bounded' in two senses: to refer to bounded sets and to refer to bounded functions. Remember that a set $A \subseteq \mathbb{R}$ is bounded if there is a real number M such that $|x| \leq M$ for each $x \in A$; and a function $f\colon X \to \mathbb{R}$ is bounded if there is a real number M with $|f(x)| \leq M$ for each $x \in X$ (so a bounded function is a function that has a bounded image set).

See Book D, Chapter 13, Definition 2.14.

When trying to identify the precise characteristics of a mathematical object that enable us to prove a result, it is often helpful to look at what happens when we vary or drop some of the conditions. The Extreme Value Theorem applies to closed and bounded intervals of $\mathbb{R}$, so we remove each of these in turn to see if both appear necessary for the result to hold.

First we examine the effect of removing the boundedness of the domain. Consider the identity function, $\mathrm{Id}\colon [0,\infty) \to \mathbb{R}$ given by $x \mapsto x$. This (rather basic) continuous function is not bounded when defined on the unbounded interval $[0,\infty)$. This suggests that the bounded condition on the domain is a necessary one. Now consider the function $f\colon [1,2) \to \mathbb{R}$, given by $x \mapsto x^2$. The domain, $[1,2)$, is bounded but is not $d^{(1)}$-closed (when considered as a subset of $(\mathbb{R}, d^{(1)})$). The function is still bounded (as $1 \leq f(x) \leq 4$), but it does not *attain* its upper bound (as there is no value $a \in [1,2)$ with $f(a) = 4$). Therefore the domain of a real function needs to be both bounded and closed in order for the function to be bounded and to attain its bounds. So it appears reasonable to start our quest for a generalisation of the Extreme Value Theorem by examining the nature of closed and bounded subsets of $\mathbb{R}$.

Here we could, of course, have used the identity function (restricted to an appropriate interval).

1.1 Preliminaries: sequences and subsequences

We use sequences and subsequences as our main tools for investigating the nature of closed and bounded subsets of $\mathbb{R}$. Therefore we first remind ourselves of a few key definitions, and warm up with a few examples.

Let X be a set. A **sequence** in X is an unending ordered list of elements of X

$$a_1, a_2, a_3, \ldots.$$

The element a_k is the ***k*th term** of the sequence, and the whole sequence is denoted by (a_k), $(a_k)_{k=1}^{\infty}$ or $(a_k)_{k\in\mathbb{N}}$.

A real sequence (a_k) is **monotonic increasing** if for each $k \in \mathbb{N}$, $a_k \leq a_{k+1}$ and **monotonic decreasing** if for each $k \in \mathbb{N}$, $a_k \geq a_{k+1}$. It is **strictly monotonic increasing or decreasing** if the inequality is strict for each k.

The sequence $(a_{n_k})_{k=1}^{\infty}$ is a **subsequence** of the sequence (a_n) if $(n_k)_{k=1}^{\infty}$ is a strictly increasing sequence of positive integers, i.e. if $1 \leq n_1 < n_2 < n_3 < \cdots$; see Figure 1.1.

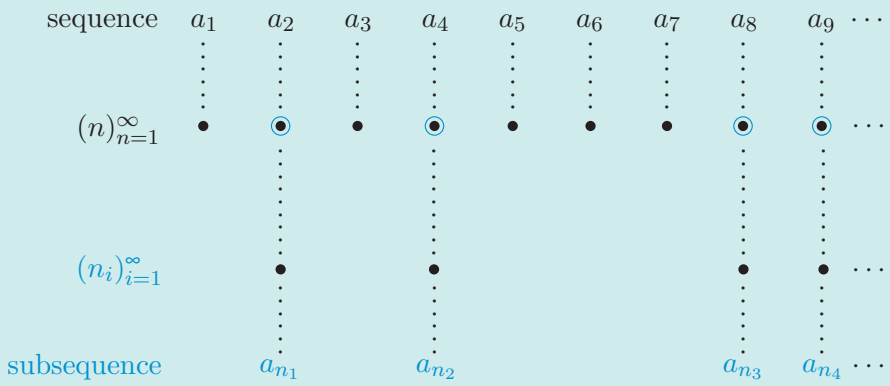

Figure 1.1

A set $A \subset \mathbb{R}$ is **bounded** if there exists a real number M such that $|x| \leq M$ for each $x \in A$.

Exercise 1.1

It is not possible to define a strictly monotonic increasing sequence for an *arbitrary* metric space by saying 'a sequence (a_n) is **strictly monotonic increasing** if for each $n \in \mathbb{N}$, $a_n < a_{n+1}$.' Why is this?

Exercise 1.2

Write down the first few terms of a sequence in $C[0, 1]$.

Exercise 1.3

See Book D, Chapter 15, Section 2.

Write down the first few terms of a sequence in Cantor space.

Exercise 1.4

See Book D, Chapter 16, Example 1.19.

Write down a null sequence in the middle-third Cantor set which does not contain the point 0 as a term.

We now focus on a concrete example and consider the set $[0, 1]$. It is very easy to write down a sequence which is not $d^{(1)}$-convergent: any alternating sequence, such as $(0, 1, 0, 1, 0, \ldots)$, will suffice. However, this sequence does have a convergent *sub*sequence. (In fact, it has infinitely many different convergent subsequences.) Since the interval $[0, 1]$ contains uncountably many points from which we can pick the terms for our sequence, and since we only need to select countably many of these to form our sequence, it may seem as if it should be easy to find a sequence that has no convergent subsequence. However, this turns out to be impossible!

See Definition 6.4 of Chapter 16.

In fact, we prove that it is impossible to select a sequence of points *from any $d^{(1)}$-closed and bounded interval* $[a, b]$ that does not have a $d^{(1)}$-convergent subsequence. In other words, we cannot pick countably many points (to form a sequence) from our uncountable set $[a, b]$, without returning arbitrarily close to some fixed point (the limit of the convergent subsequence) infinitely many times.

Of course this is not true for all metric spaces; the sequence (n) in $\mathbb{R}$ clearly has no convergent subsequence. However, we do not have to look to the whole of $\mathbb{R}$ to find an example. Consider the bounded, open interval $(0, 1)$. The sequence $\left(\frac{1}{n}\right)$ has no $d^{(1)}_{(0,1)}$-convergent subsequence. In the case of (n) in $\mathbb{R}$, we use the fact that $\mathbb{R}$ is unbounded to give us enough 'room' to find a sequence with no convergent subsequence. In the case of $\left(\frac{1}{n}\right)$ in $(0, 1)$, we use the fact that $(0, 1)$ has a closure point outside the interval in order to find 'room'. If you lived inside the interval $(0, 1)$, then you would find that the edges appeared 'fuzzy' – you could approach arbitrarily close to the points 0 or 1, but you could not reach either of them. In fact, the 'fuzzy' edges of $(0, 1)$ may look a little like the unreachable 'edges' of $\mathbb{R}$ (and you may recall that in Chapter 21 we proved that $(0, 1)$, with the Euclidean metric, was homeomorphic to $\mathbb{R}$, with the Euclidean metric).

This works because we can find a sequence of points that converges to a closure point; see Book D, Chapter 16, Worked Exercise 3.7.

Our choice of $[0, 1]$ as our closed and bounded interval was fairly random. It turns out that we can always find a convergent subsequence for any sequence in any $d^{(1)}$-closed and bounded subset of $\mathbb{R}$! As the first step on our journey towards a proof, we prove a lemma.

Lemma 1.1

Let (a_n) be a sequence of real numbers. Then (a_n) contains a monotonic subsequence – that is, there is a subsequence (a_{n_k}) such that either:

$$a_{n_1} \leq a_{n_2} \leq a_{n_3} \leq \cdots \leq a_{n_k} \leq \cdots$$

or

$$a_{n_1} \geq a_{n_2} \geq a_{n_3} \geq \cdots \geq a_{n_k} \geq \cdots .$$

The lemma does not state that the monotonic subsequence converges: the sequence $(1, -1, 2, -2, 3, -3, 4, -4, \ldots)$ contains the strictly increasing subsequence $(1, 2, 3, \ldots)$, which does not converge.

Proof Let (a_n) be a sequence in $\mathbb{R}$. We call a particular term a_i of (a_n) a *peak point* if $a_i \geq a_m$ for all $m \geq i$ (Figure 1.2). If (a_n) has infinitely many peak points, then the sequence of peak points $\left(a_{n_i}\right)$, where $n_1 < n_2, \ldots$, is a monotonic decreasing subsequence of (a_n) – and we are done.

If this isn't the case, then there are only finitely many peak points in (a_n). Let $N_1 = n_i + 1$, where a_{n_i} is the final peak point (setting $N_1 = 1$ if (a_n) has no peak points). Since a_{N_1} is not a peak point, then there must be a point a_{N_2}, further along the sequence (so $N_2 > N_1$) with $a_{N_2} > a_{N_1}$. In Exercise 1.5 we ask you to use induction to finish constructing a monotonic increasing subsequence $\left(a_{N_i}\right)$ of (a_n).

Therefore (a_n) contains a monotonic subsequence and, as (a_n) was an arbitrary sequence in $\mathbb{R}$, we have the result. ■

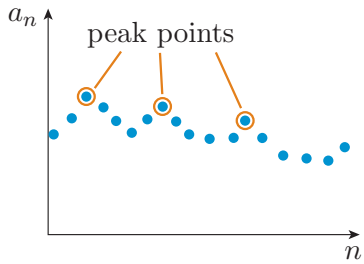

Figure 1.2

Exercise 1.5

Write a formal proof using induction to show that if there are only finitely many peak points, then we can construct a monotonic increasing subsequence (a_{N_i}) of the sequence (a_n) in the proof of Lemma 1.1.

Theorem 1.2

Let X be any $d^{(1)}$-closed and bounded subset of $\mathbb{R}$. Then every sequence in X contains a $d_X^{(1)}$-convergent subsequence.

Proof By Lemma 1.1 we know that every sequence contains a monotonic subsequence. If this is a monotonic increasing sequence then, since X is bounded above, we can apply the Monotone Convergence Theorem and conclude that this subsequence is $d^{(1)}$-convergent when considered as a sequence in $\mathbb{R}$. As X is $d^{(1)}$-closed, then the limit of this sequence must be in X. Similarly, if this is a monotonic decreasing sequence then, since X is bounded below, we can also apply the Monotone Convergence Theorem and draw the same conclusion, giving the result. ■

The Monotone Convergence Theorem is Theorem 1.22 of Chapter 13.

We now turn our attention back to the Boundedness Theorem and the Extreme Value Theorem and prove these results for $\mathbb{R}$.

Theorem 1.3 *Boundedness Theorem*

Let $f\colon [a,b] \to \mathbb{R}$ be a $(d_{[a,b]}^{(1)}, d^{(1)})$-continuous function on $[a,b]$. Then f is bounded on $[a,b]$. That is, there is a real number M such that $|f(x)| \leq M$ for each $x \in [a,b]$.

Proof The proof is by contradiction. Supposed that f is not bounded. Then for each $n \in \mathbb{N}$ we can find a point $x_n \in [a,b]$ with $|f(x_n)| > n$. This gives us two sequences. One, (x_n), consisting of points in $[a,b]$, and the second, $\big(f(x_n)\big)$, consisting of points in $\mathbb{R}$.

First consider the second sequence, $\big(f(x_n)\big)$. This sequence was chosen so that $|f(x_n)| > n$ for each n. Therefore $\frac{1}{f(x_n)}$ is a null sequence by the Squeeze Rule (since $|\frac{1}{f(x_n)}| < \frac{1}{n}$ for each n and $\big(\frac{1}{n}\big)$ is a basic null sequence). Therefore $\big(f(x_n)\big)_{n=1}^{\infty}$ is divergent (by the remarks after Book D, Chapter 13, Worked Exercise 1.18). Moreover, we can apply the same reasoning to every subsequence $\big(f(x_{n_i})\big)$ of $\big(f(x_n)\big)$ (since $f(x_{n_i}) = f(x_j)$ for some $j \geq i$ and so $\big|\frac{1}{f(x_{n_i})}\big| \leq |\frac{1}{i}|$).

The Squeeze Rule is Lemma 1.15 of Chapter 13.

Now consider the first sequence, (x_n). Since this is a sequence of points in a closed interval, by Theorem 1.2 this sequence has a convergent subsequence, say (x_{n_i}), that is $d^{(1)}$-convergent to a point, say x, in the interval $[a,b]$. (We know that $x \in [a,b]$ by Book D, Chapter 16, Worked

Exercise 3.7.) Now f is continuous and so the limit of the sequence $(f(x_{n_i}))$ is $f(x)$, contradicting the fact that $(f(x_{n_i}))$ is not convergent. Therefore f is bounded and we have the result. ■

Theorem 1.4 *Extreme Value Theorem*

Let $f\colon [a,b] \to \mathbb{R}$ be a $(d^{(1)}_{[a,b]}, d^{(1)})$-continuous function on $[a,b]$. Then there are numbers $c, d \in [a,b]$ such that

$$f(c) \leq f(x) \leq f(d) \quad \text{for each } x \in [a,b].$$

Proof Let $f\colon [a,b] \to \mathbb{R}$ be a continuous function on $[a,b]$. Then by Theorem 1.3 there is a real number M such that $|f(x)| \leq M$ for each $x \in [a,b]$. Therefore

$$-M \leq f(x) \leq M$$

for each $x \in [a,b]$.

We know that the image set $A = \{y \in \mathbb{R}\colon y = f(x) \text{ for some } x \in [a,b]\}$ is bounded above (by M) and so by the least upper bound property of the real numbers, the set A has a least upper bound, say L. Therefore we need to find a $d \in [a,b]$ with $f(d) = L$. We do this by constructing a sequence of points in $[a,b]$, (a_n), such that $f(a_n)$ gets closer and closer to L. We then use the fact that this sequence has a convergent subsequence to finish the proof.

See least upper bound property in the *Handbook*.

Since L is a least upper bound for the image set A, we have that for each n, $L - \frac{1}{n}$ is not an upper bound for A. Therefore we can find a point $a_n \in [a,b]$ with $L - \frac{1}{n} \leq f(a_n) \leq L$. This defines a sequence (a_n) in $[a,b]$. Therefore we now have two sequences: (a_n) in $[a,b]$, and $(f(a_n))$ in $A \subseteq \mathbb{R}$.

The sequence $(f(a_n))$ is $d^{(1)}$-convergent to L (when considered as a sequence in $\mathbb{R}$), but it may not be $d^{(1)}_A$-convergent (as L may not be in A).

For example, the sequence $\left(\frac{1}{n}\right)$, as a sequence in $\mathbb{R}$, is $d^{(1)}$-convergent to 0; but when considered as a sequence in $(0,1]$, $\left(\frac{1}{n}\right)$ is not $d^{(1)}_{(0,1]}$-convergent.

Now $[a,b]$ is a $d^{(1)}_{[a,b]}$-closed and bounded subset of $\mathbb{R}$, and so by Theorem 1.2, (a_n) has a convergent subsequence $\left(a_{n_i}\right)$. Set $d \in [a,b]$ equal to the limit of this convergent subsequence. (We know that $d \in [a,b]$ by Book D, Chapter 16, Worked Exercise 3.7.) Now since f is continuous, by Book D, Chapter 13, Definition 2.1, we know that the subsequence $f(a_{n_i}) \to f(d)$. Since every subsequence of a convergent sequence converges to the limit of the sequence (by Book D, Chapter 16, Lemma 1.13), we have $f(d) = L$ and f attains its upper bound.

Here we use the existence of a convergent subsequence for any sequence in $[a,b]$.

A similar argument shows that f attains its lower bound. ■

Remarks

In the proof of the Boundedness Theorem, we assumed that the function, f, was not bounded and used this assumption to construct a sequence in a closed and bounded interval. We then used the fact that this sequence had a convergent subsequence to get our contradiction.

In the proof of the Extreme Value Theorem, we used the Boundedness Theorem to find a bound on the function and then used the existence of this bound to construct a sequence in a closed and bounded interval. Again we used the fact that this sequence had a convergent subsequence to prove that the function attained its bounds. Therefore the important characteristic of the closed and bounded interval seems to be the fact that every sequence has a convergent subsequence.

2 Sequential compactness

In the proof of the Extreme Value Theorem, the key property of the $d^{(1)}$-closed and bounded interval $[a, b]$ that allows us to show that the function attains its bounds is the fact that the sequence (a_n) has a convergent subsequence. It is this property that we use to define *sequential compactness*, a generalisation of '$d^{(1)}$-closed and bounded'-ness in $\mathbb{R}$.

Definition 2.1 *d-sequentially compact*

A metric space (X, d) is **d-sequentially compact** if each sequence in X has a d-convergent subsequence (which is d-convergent to a point in X).

A subset $A \subseteq X$ is a **sequentially compact subset** if (A, d_A) is sequentially compact as a metric space.

Remarks

1. If the metric is clear from the context, then we suppress mention of the metric and say that X is *sequentially compact*.
2. If $A \subseteq X$ is a sequentially compact subset, then we sometimes refer to A as d_A-sequentially compact.
3. It is an essential part of this definition that the convergent subsequence converges to a point in the set under consideration.

Exercise 2.1

Prove that any sequentially compact subset of a metric space (X, d) is d-closed.

We wanted to generalise '$d^{(1)}$-closed and bounded'-ness in $\mathbb{R}$, and Exercise 2.1 shows that we have achieved part of this objective. The next result again links sequential compactness with '$d^{(1)}$-closed and bounded'-ness.

Corollary 2.2

Any $d^{(1)}$-closed and bounded subset of $\mathbb{R}$ is sequentially compact.

Proof This is an immediate corollary of Theorem 1.2. ■

In fact, it turns out that the sequentially compact subspaces of $\mathbb{R}$ are precisely those which are $d^{(1)}$-closed and bounded.

Theorem 2.3 *The Heine–Borel Theorem for* $\mathbb{R}$

A subset of $\mathbb{R}$ is $d^{(1)}$-sequentially compact if, and only if, it is $d^{(1)}$-closed and bounded.

Proof Corollary 2.2 states that any $d^{(1)}$-closed and bounded subset of $\mathbb{R}$ is sequentially compact. We need to prove that any sequentially compact subset of $\mathbb{R}$ (with the Euclidean metric) is both $d^{(1)}$-closed and bounded.

By Exercise 2.1, a $d^{(1)}$-sequentially compact subset of $\mathbb{R}$ is $d^{(1)}$-closed. The proof that A is bounded is by contradiction. Assuming that A is not bounded, it is possible to construct a sequence in A which does not have a convergent subsequence. We leave this for you as Exercise 2.2. Therefore, by contradiction, A is bounded and we have the result. ■

Exercise 2.2

Show that every sequentially compact subset, A, of $\mathbb{R}$ (with the usual Euclidean metric), is bounded by assuming that A is not bounded and constructing a sequence in A that does not have a convergent subsequence.

2.1 Examples of sequentially compact spaces

We know, from Theorem 2.3 that the sequentially compact subsets of $\mathbb{R}$ are precisely the $d^{(1)}$-closed and bounded subsets of $\mathbb{R}$. We now look at some other examples and non-examples.

Exercise 2.3

Is K_C, the middle-third Cantor set (with the metric induced by the Euclidean metric on $\mathbb{R}$), sequentially compact? Prove your answer.

The middle-third Cantor set is defined in Book D, Chapter 16, Example 1.19.

Remember that to check that a set is sequentially compact you must show that *every* sequence has a convergent subsequence.

Exercise 2.4

Let (X, d) be a metric space where X is a finite set and d is any metric. Show that X is sequentially compact.

To show that a set is not sequentially compact just requires finding *one* sequence without a convergent subsequence.

Worked Exercise 2.4

Let

$$A = \{f \in C[0,1] : |f(x)| \leq 1 \text{ for all } x \in [0,1]\}.$$

Show that A is not sequentially compact for $d_{\max}$.

See Book D, Chapter 15, Definition 5.1.

Solution

We need to find just one sequence of functions in the set A that has no subsequence converging to a point in the set A.

With this in mind, we consider the sequence (f_n) in A given by

$$f_n(x) = \begin{cases} 1 - nx, & \text{if } 0 \leq x < \frac{1}{n}, \\ 0, & \text{if } \frac{1}{n} \leq x \leq 1. \end{cases}$$

The first few functions in the sequence are shown in Figure 2.1.

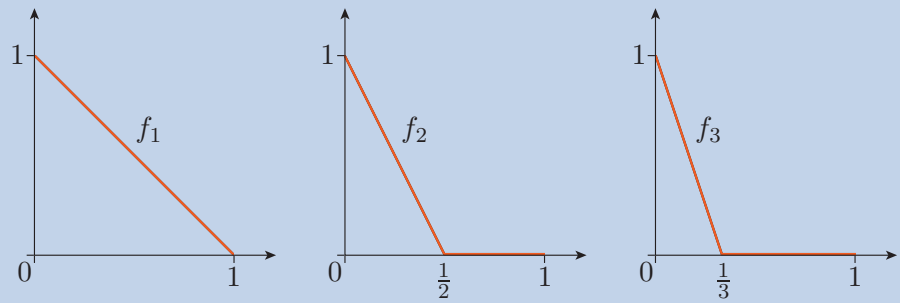

Figure 2.1

We know that, for $x \in [0,1]$, the limit as n tends to infinity of $f_n(x)$ is given by:

$$f_n(x) \to \begin{cases} 1, & \text{if } x = 0, \\ 0, & \text{if } 0 < x \leq 1. \end{cases}$$

Thus the pointwise limit of the sequence is a discontinuous function. Moreover, any subsequence of (f_n) has the same pointwise limit. Hence no subsequence can converge to a continuous function and A is not sequentially compact.

We investigated this sequence of functions and calculated the limit in Book D, Chapter 15, Worked Exercise 5.9.

Exercise 2.5

Let

$$A = \{f \in C[0,1] : |f(x) - f(y)| \leq |x - y| \text{ for all } x, y \in [0,1]\}.$$

Show that A is not sequentially compact for $d_{\max}$.

Hint: find a sequence (f_n) of functions in A such that (f_n) and any subsequence of (f_n) have no pointwise limit, and then make use of results from Book D, Chapter 15.

2.2 Subsets, images and products of sequentially compact spaces

Suppose that (X, d) is a d-sequentially compact metric space. It is natural to ask whether there are any subsets A of X which are necessarily d_A-sequentially compact. We know from Exercise 2.4 that any finite subset is sequentially compact. The next result tells us that any d-closed subset A of X is d_A-sequentially compact.

Theorem 2.5

Let $A \subseteq X$ where (X, d) is a d-sequentially compact metric space. Then A is d-closed if, and only if, A is d_A-sequentially compact.

Proof We first show that a d-closed subset is sequentially compact.

Let A be a d-closed subset of a sequentially compact metric space (X, d) and let (a_n) be a sequence in A. We need to find a subsequence of (a_n) which is d_A-convergent. Since X is sequentially compact, and (a_n) is a sequence in X, we know that (a_n) has a d-convergent subsequence, say (a_{n_i}); we need to prove that the limit of this subsequence is in A. However, since A is closed, by definition, if we have any d-convergent sequence with terms in A, then the limit must lie in A. The subsequence (a_{n_i}) is such a sequence and so the limit of this must lie in A. Therefore (a_{n_i}) is d_A-convergent and A is sequentially compact.

By Exercise 2.1, any sequentially compact subset of any metric space (X, d) is d-closed, and so sequentially compact subsets of sequentially compact metric spaces are certainly closed. ■

Remark

In particular, by taking $A = X$, we see that sequentially compact sets are closed.

Since convergence is preserved by continuous functions, it is no surprise that the continuous image of a sequentially compact set is sequentially compact.

Lemma 2.6

Let (X, d) and (Y, e) be metric spaces with X being sequentially compact. Let $f: X \to Y$ be a (d, e)-continuous function. Then the image, $f(X)$, is sequentially compact.

Proof This is Exercise 2.6. ■

Exercise 2.6

Prove Lemma 2.6.

This result tells us that sequential compactness is a topological invariant.

Corollary 2.7

Sequential compactness is preserved by homeomorphism.

Proof Let (X, d) and (Y, e) be homeomorphic metric spaces and let $f: X \to Y$ be a homeomorphism between them. The map $f: X \to Y$ is an onto and (d, e)-continuous map. Therefore if X is sequentially compact then Lemma 2.6 implies that Y must be sequentially compact. Also $f^{-1}: Y \to X$ is an onto (e, d)-continuous map from Y to X so, by the same lemma, X is sequentially compact if Y is sequentially compact. Therefore X is sequentially compact if, and only if, Y is sequentially compact. ■

We know that many properties are preserved by taking products of metric spaces. Sequential compactness is one of these properties.

Theorem 2.8 *Tikhonov's Theorem*

Let (X, d) and (Y, d') be two sequentially compact metric spaces. Then $(X \times Y, e)$ is sequentially compact (where e is a product metric).

Proof We need to show that every sequence in $X \times Y$ contains an e-convergent subsequence. Let (x_n, y_n) be a sequence in $X \times Y$. First consider (x_n) as a sequence in X. Since X is sequentially compact, this sequence has a convergent subsequence, which we write as (x_{n_i}). Now consider the sequence (y_{n_i}) in Y. This is not, of course, necessarily convergent but, since Y is sequentially compact, it also has a convergent subsequence, which we write as $(y_{n_{i_j}})$.

Now consider the sequence $(x_{n_{i_j}}, y_{n_{i_j}})$. By Book D, Chapter 15, Theorem 3.15, this sequence is e-convergent to a point (x, y) in $X \times Y$ if the component sequence $\left(x_{n_{i_j}}\right)$ is d-convergent to x in X and the component sequence $\left(y_{n_{i_j}}\right)$ is d'-convergent to y in Y. We know that

$(y_{n_{i_j}})$ is convergent to a point in Y and $(x_{n_{i_j}})$ is a subsequence of the convergent sequence (x_{n_i}), and so converges to the same point in X (by Book D, Chapter 16, Lemma 1.13) – see Figure 2.2.

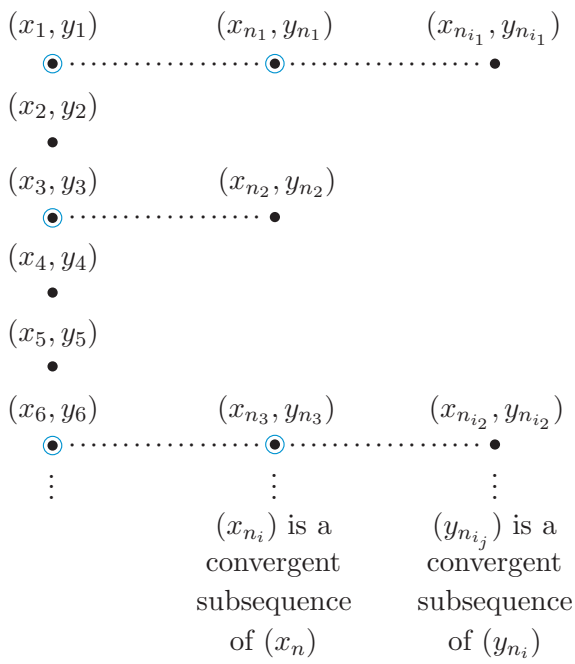

Figure 2.2

Therefore $(x_{n_{i_j}}, y_{n_{i_j}})$ is a convergent subsequence of our arbitrary sequence (x_n, y_n) and $X \times Y$ is sequentially compact. ■

2.3 Sequentially compact subsets of $\mathbb{R}^n$

We know from Corollary 2.2 that any $d^{(1)}$-closed and bounded subset of $\mathbb{R}$ (with the usual Euclidean metric) is sequentially compact. We can use Theorem 2.8 to show that any $d^{(n)}$-closed and bounded subset of $\mathbb{R}^n$ is sequentially compact, but first we need to say what we mean by the term *bounded* when we are in a metric space other than $(\mathbb{R}, d^{(1)})$.

Boundedness is defined for subsets of $\mathbb{R}$ (with the Euclidean metric) in Book D, Chapter 13, Definition 2.14: a set $A \subset \mathbb{R}$ is *bounded* if there exists a real number M such that $|x| \leq M$ for each $x \in A$.

Definition 2.9 *Bounded set*

Let (X, d) be a metric space and let $A \subseteq X$. Then A is **bounded** if there is $M > 0$ such that $d(x, y) \leq M$ for all $x, y \in A$. If we want to emphasise the metric, we say that A is **d-bounded**.

Intuitively, a subset A of a metric space is bounded if we can find a number M such that all the points in A are within M of one another. This number M can be as large as we like, but it must be finite.

Worked Exercise 2.10

Show that the definition of a bounded subset of $\mathbb{R}$ given in Book D, Chapter 13, Definition 2.14 is equivalent to Definition 2.9 for the Euclidean metric on $\mathbb{R}$.

Solution

Let $A \subseteq \mathbb{R}$ be bounded. If $A = \emptyset$, then there is nothing to prove, so suppose $A \neq \emptyset$. Suppose that A is bounded as defined in Definition 2.9. Then there exists $M > 0$ such that $d^{(1)}(x, y) \leq M$ for all $x, y \in A$. Hence, fixing some $y \in A$ and using the Triangle Inequality, we find for any $x \in A$,

$$|x| = |x - y + y| \leq |x - y| + |y| \leq M + |y|.$$

Thus A is bounded according to Definition 2.14 of Chapter 13.

Now suppose that there exists $M > 0$ such that $|x| \leq M$ for all $x \in A$. Then, by the Triangle Inequality,

$$|x - y| \leq |x| + |y| \leq M + M = 2M,$$

for all $x, y \in A$. Thus A is bounded as defined in Definition 2.9.

Exercise 2.7

The projection functions are defined in Book D, Chapter 14, Definition 1.18.

Show that if $A \subseteq \mathbb{R}^n$ is $d^{(n)}$-bounded, and p_i is the ith projection function then the image of A under this function, $p_i(A) \in \mathbb{R}$ is $d^{(1)}$-bounded.

The definition of bounded for a subset A of a metric space (X, d) is equivalent to saying that we can enclose A within a ball (as we ask you to show in the next exercise).

Exercise 2.8

Show that the definition of bounded as given in Definition 2.9 is equivalent to the statement that A can be enclosed within an open ball $B_d(a, r)$ for some $a \in A$ and $r > 0$.

For $\mathbb{R}^n$ we can show that any closed and bounded set is $d^{(n)}$-sequentially compact.

Proposition 2.11

Any $d^{(n)}$-closed and bounded subset of $\mathbb{R}^n$ is $d^{(n)}$-sequentially compact.

Proof Suppose that $K \subseteq \mathbb{R}^n$ is $d^{(n)}$-closed and bounded. Since K is bounded, by Exercise 2.7 $p_j(K)$ is bounded for each $j = 1, 2, \ldots, n$, where p_j is the usual projection function onto the jth coordinate axis. So we can find $M_j > 0$ such that

$$p_j(K) \subseteq [-M_j, M_j], \quad j = 1, 2, \ldots, n.$$

Hence

$$K \subseteq [-M_1, M_1] \times [-M_2, M_2] \times \cdots \times [-M_n, M_n]$$

(Figure 2.3). Each of the intervals $[-M_j, M_j]$ is a sequentially compact subset of $\mathbb{R}$. Thus, by Tikhonov's Theorem, the product of the intervals is also sequentially compact. Hence K is a closed subset of a sequentially compact set, and so is sequentially compact, by Theorem 2.5. ■

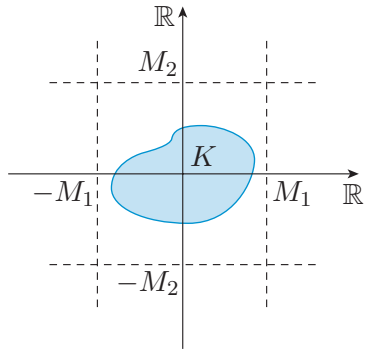

Figure 2.3

Exercise 2.9

Show that S^3, the unit sphere in $\mathbb{R}^4$, given by $\{(w, x, y, z) \in \mathbb{R}^4 : w^2 + x^2 + y^2 + z^2 = 1\}$, is sequentially compact (for the Euclidean metric). (*Hint*: you may find Book D, Chapter 16, Theorem 5.2 helpful.)

We now know that all $d^{(n)}$-closed and bounded subsets of $\mathbb{R}^n$ are sequentially compact. We next show that these are the only sequentially compact subsets of $\mathbb{R}^n$.

Theorem 2.12 *The Heine–Borel Theorem*

A subset of $\mathbb{R}^n$ is $d^{(n)}$-sequentially compact if, and only if, it is $d^{(n)}$-closed and bounded.

Proof We know from Proposition 2.11 that any $d^{(n)}$-closed and bounded subset of $\mathbb{R}^n$ is sequentially compact. Therefore it is sufficient to prove that any $d^{(n)}$-sequentially compact subset of $\mathbb{R}^n$ is $d^{(n)}$-closed and bounded.

We know from Exercise 2.1 that a sequentially compact subset of any metric space (X, d) is d-closed, so this is certainly true for a subset of $\mathbb{R}^n$. Therefore it is sufficient to prove that A is bounded.

The proof that A is bounded is by contradiction. Assume that A is not bounded. Then A cannot be empty and so we may assume that there is an element $a_1 \in A$. By Exercise 2.8 we cannot find an $M > 0$ such that $B_{d^{(n)}}(a_1, M) \supseteq A$. Therefore for each $n \in \mathbb{N}$ we can find an $a_n \in A$ with $d(a_1, a_n) > n$. This gives us a sequence (a_n) in A.

Since $a_n \in A$, and A is sequentially compact, (a_n) has a convergent subsequence (a_{n_i}). Suppose that a is the limit of (a_{n_i}). Then we can find

$I \in \mathbb{N}$ so that for $i > I$, $d(a_{n_i}, a) < 1$. Thus for $i > I$, on using the Reverse Triangle Inequality for d, we find that

$$\begin{aligned} d(a_1, a) &\geq |d(a_1, a_{n_i}) - d(a_{n_i}, a)| \\ &\geq d(a_1, a_{n_i}) - 1 \\ &\geq n_i - 1, \text{ for each } i > I. \end{aligned}$$

But $n_i \to \infty$ and so $d(a_1, a)$ is not finite, which is impossible; we conclude that A is bounded. ■

As promised we now generalise the Extreme Value Theorem.

2.4 The Generalised Extreme Value Theorem

We can now prove a general version of the Extreme Value Theorem.

Theorem 2.13 *General Extreme Value Theorem*

Let (X, d) be a non-empty sequentially compact metric space, and let $f\colon X \to \mathbb{R}$ be $(d, d^{(1)})$-continuous. Then there are $c, d \in X$ such that

$$f(c) \leq f(x) \leq f(d) \quad \text{for all } x \in X.$$

Proof Since X is sequentially compact and f is continuous, Lemma 2.6 implies that $f(X)$ is a sequentially compact subset of $\mathbb{R}$. Hence, by Theorem 2.3, $f(X)$ is a closed and bounded subset of $\mathbb{R}$.

See upper bound and lower bound in the *Handbook*.

Since the image $f(X)$ is bounded, it has an upper bound and a lower bound. Therefore, by the least upper bound property of the real numbers, the set $f(X)$ has a least upper bound M and a greatest lower bound m. For any $r > 0$, since $M - \frac{r}{2}$ is not an upper bound for $f(X)$, we can find a point $p \in f(X)$ with $M - \frac{r}{2} \leq f(p) \leq M$. Therefore we can deduce that $B_{d^{(1)}}(M, r)$ must intersect $f(X)$ for any $r > 0$ and so, by Book D, Chapter 16, Worked Exercise 3.7, we see that M is a closure point of $f(X)$. But $f(X)$ is closed, and so we deduce that $M \in f(X)$. A similar argument shows that $m \in f(X)$. Thus there are $c \in X$ and $d \in X$ with $f(c) = m$ and $f(d) = M$. It follows from the definition of m and M that $f(c) \leq f(x) \leq f(d)$ for all $x \in X$. ■

It follows from Theorem 2.13 that any $(d^{(n)}, d^{(1)})$-continuous function from a sequentially compact subset of $\mathbb{R}^n$ to $\mathbb{R}$ has a least upper bound and a greatest lower bound, and attains them both; in other words, it has a maximum value and a minimum value.

If the function is $(d^{(n)}, d^{(1)})$-continuous but the domain is not sequentially compact, then this need not be the case.

Worked Exercise 2.14

Let $f: \mathbb{R}^2 \to \mathbb{R}$ be given by

$$f(x, y) = \sqrt{x^2 + y^2}.$$

Show that f attains a maximum value on $B_{d^{(2)}}[(0, 0), 1]$; see Figure 2.4.

Solution

In order to use Theorem 2.13 to deduce the existence of a maximum value, we must verify that f is $(d^{(2)}, d^{(1)})$-continuous on its domain and that $B_{d^{(2)}}[(0, 0), 1]$ is a sequentially compact set.

We saw in Book D, Chapter 13, Worked Exercise 3.17 that f is continuous on the whole of $\mathbb{R}^2$, and so, by the Restriction Rule, is continuous on $B_{d^{(2)}}[(0, 0), 1]$.

The ball $B_{d^{(2)}}[(0, 0), 1]$ is closed and bounded. Hence, by Theorem 2.12, it is sequentially compact.

Thus both conditions of the General Extreme Value Theorem are satisfied, and we conclude that the function f attains a maximum value on $B_{d^{(2)}}[(0, 0), 1]$.

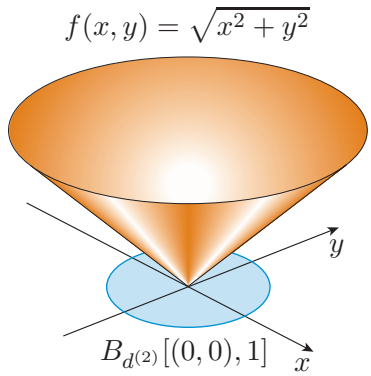

Figure 2.4

Exercise 2.10

Let $f: \mathbb{R}^2 \to \mathbb{R}$ be given by

$$f(x, y) = y \sin(\pi x) + x \cos(\pi y).$$

Show that f attains a minimum value on the set $\{(x, y) : 0 \leq x \leq 1, \text{ and } 0 \leq y \leq 1\}$.

The General Extreme Value Theorem is a powerful result since it guarantees the existence of a maximum value and a minimum value of a continuous function on a sequentially compact domain without our having to know where these values are.

We now turn our attention to subsets of the function space $C[0, 1]$.

3 Functions and sequential compactness

In this section we investigate the relationship between sequential compactness and functions in $(C[0, 1], d_{\max})$, with the goal of finding sufficient conditions to ensure that a given subset of $C[0, 1]$ is sequentially compact.

We begin by examining in a little more detail what it means for a function from $[0,1]$ to $\mathbb{R}$ to be continuous. We discover that the fact that $[0,1]$ is sequentially compact means that f satisfies a stronger notion of continuity, known as *uniform continuity*. This will lead us to identifying one of the conditions needed to guarantee sequential compactness of a subset of $C[0,1]$.

3.1 Sequentially compact subsets of $\mathbb{R}$ and continuous functions

We start by finding an alternative formulation of what it means for a function from a subset of $\mathbb{R}$ to $\mathbb{R}$ to be continuous (for the Euclidean metrics).

Proposition 3.1

Let $A \subseteq \mathbb{R}$ and $f\colon A \to \mathbb{R}$. Then f is (d_A, d)-continuous at $a \in A$ if, and only if, for each $\varepsilon > 0$, there is $\delta > 0$ such that

$$|f(x) - f(a)| < \varepsilon, \text{ whenever } x \in A \text{ and } |x - a| < \delta.$$

(Here d denotes the Euclidean metric for $\mathbb{R}$ and d_A denotes the corresponding induced metric for A.)

Proof *Only if*

Let $a \in A$ and suppose that f is (d_A, d)-continuous at $a \in A$.

Let $\varepsilon > 0$ be given. We suppose, for a contradiction, that for each $\delta > 0$, we can find $x \in B_{d_A}(a, \delta)$ for which $|f(x) - f(a)| \geq \varepsilon$. Then, in particular, for each $n \in \mathbb{N}$, we can find $x_n \in B_{d_A}(a, \frac{1}{n})$ for which $|f(x_n) - f(a)| \geq \varepsilon$. Thus for this sequence (x_n), we have that it converges to a and yet $|f(x_n) - f(a)| \geq \varepsilon > 0$ for each $n \in \mathbb{N}$, meaning that the sequence $(f(x_n))$ does not converge to $f(a)$ – this contradicts the assumption that f is continuous at a.

If

For the opposite implication, let $a \in A$ and suppose that for each $\varepsilon > 0$, there is $\delta > 0$ such that $|f(x) - f(a)| < \varepsilon$, whenever $x \in A$ and $|x - a| < \delta$.

Let (x_n) be a sequence in A that d_A-converges to a. Our aim is to show that $f(x_n)$ converges to $f(a)$ as $n \to \infty$.

Let $\varepsilon > 0$ be given. Then, by our assumption, we can find $\delta > 0$ for which $|f(x) - f(a)| < \varepsilon$, whenever $d_A(x, a) < \delta$. However, since $d_A(x_n, a) \to 0$ as $n \to \infty$, we can find $N \in \mathbb{N}$ so that $d_A(x_n, a) < \delta$ for $n > N$, and so for $n > N$, $d(f(x_n), f(a)) < \varepsilon$. In other words, for each $\varepsilon > 0$, there is $N \in \mathbb{N}$ so that $d(f(x_n), f(a)) < \varepsilon$ for $n > N$, which means that $f(x_n)$ converges to $f(a)$ as $n \to \infty$, and we conclude that f is continuous at a, as required. ■

Remark

This equivalent formulation of continuity is known as the ε-δ **definition of continuity** – you will likely have met it in your earlier studies. It is discussed in M208.

An examination of the proof reveals that there is no particular requirement to restrict our attention to a subset of the real line, and a similar result can be proved in general metric spaces.

Theorem 3.2

Let (X, d) and (Y, e) be metric spaces and let $f\colon X \to Y$. Then f is (d, e)-continuous at $a \in X$ if, and only if, for each $\varepsilon > 0$, there is $\delta > 0$ so that $e(f(x), f(a)) < \varepsilon$, whenever $d(x, a) < \delta$.

Proof See Exercise 3.1. ■

Exercise 3.1

Prove Theorem 3.2.

These results are only about continuity and we have not made use of sequential compactness – in our next results, we look at what happens in the case that the domain of the function is sequentially compact (and so, in particular, there is a metric available). Our ultimate goal in this section is to understand what it means for a subset of $C[0, 1]$ to be sequentially compact. We know from Corollary 2.2 that the interval $[0, 1]$ (and any other closed bounded interval $[a, b]$) is itself a sequentially compact subset of $\mathbb{R}$ for the Euclidean metric, and so we start by looking at a single continuous function with this domain.

Theorem 3.3

Let $f\colon [0, 1] \to \mathbb{R}$ be a $(d^{(1)}_{[0,1]}, d^{(1)})$-continuous function. Then for each $\varepsilon > 0$, there is $\delta > 0$ such that for each $x, y \in [0, 1]$, $|f(x) - f(y)| < \varepsilon$, whenever $|x - y| < \delta$.

Remarks

1. You should contrast this with the statement of Proposition 3.1 to see how these results differ: in Proposition 3.1, the choice of δ will in general depend on both ε and the point a, whereas here the same δ will do for *all* points x in $[0, 1]$.
2. A similar result holds if the interval $[0, 1]$ is replaced by a closed bounded interval $[a, b]$.

Proof We prove the theorem using proof by contradiction. Suppose that there is $\varepsilon > 0$ such that for each $\delta > 0$ there are $x, y \in [0,1]$ for which $|x - y| < \delta$ and yet $|f(x) - f(y)| \geq \varepsilon$.

Then, in particular, for each $n \in \mathbb{N}$, we can find $x_n, y_n \in [0,1]$ for which $|x_n - y_n| < \frac{1}{n}$ and yet $|f(x_n) - f(y_n)| \geq \varepsilon$.

Consider the sequence (x_n). Since each term lies in the sequentially compact set $[0,1]$, there is a convergent subsequence $(x_{n_i})_{i \in \mathbb{N}}$, with limit x, say, that lies in $[0,1]$. Since f is continuous, we know that $f(x_{n_i}) \to f(x)$ as $n \to \infty$.

Let us look at the corresponding subsequence (y_{n_i}) of (y_n). Then we have $|x_{n_i} - y_{n_i}| < \frac{1}{n_i} \leq \frac{1}{i}$ for each $i \in \mathbb{N}$, and so the Triangle Inequality for modulus implies that

$$|x - y_{n_i}| \leq |x - x_{n_i}| + |x_{n_i} - y_{n_i}| \leq |x - x_{n_i}| + \frac{1}{i}.$$

But $|x - x_{n_i}| \to 0$ as $i \to \infty$ and so, by the Squeeze Rule, we deduce that $y_{n_i} \to x$ as $i \to \infty$. Hence, since f is continuous, $f(y_{n_i}) \to f(x)$ as $i \to \infty$.

However, for each $i \in \mathbb{N}$, the definition of x_{n_i} and y_{n_i}, together with the Triangle Inequality for modulus, implies that

$$\begin{aligned} 0 < \varepsilon &\leq |f(x_{n_i}) - f(y_{n_i})| \\ &\leq |f(x_{n_i}) - f(x)| + |f(x) - f(y_{n_i})| \to 0 \text{ as } i \to \infty, \end{aligned}$$

which is a contradiction. We conclude that the theorem holds. ■

If you look closely at the proof of this result, you will see that the only property about $[0,1]$ that we made use of is that it is sequentially compact: an almost identical argument shows that a similar result holds in the case when the domain of the continuous function is sequentially compact.

Theorem 3.4

Let (X, d) and (Y, e) be metric spaces and suppose $f : X \to Y$ is (d, e)-continuous on X.

If X is d-sequentially compact, then for each $\varepsilon > 0$, there is $\delta > 0$ such that for each $x, y \in X$, $e(f(x), f(y)) < \varepsilon$ whenever $d(x, y) < \delta$.

Proof See Exercise 3.2. ■

Exercise 3.2

Prove Theorem 3.4.

This property (of continuous functions on sequentially compact sets) is important enough to have been given a name: *uniform continuity*.

Definition 3.5 *Uniform continuity*

Let (X, d) and (Y, e) be metric spaces and suppose that $f\colon X \to Y$ satisfies the following condition: for each $\varepsilon > 0$, there is $\delta > 0$ such that for each $x, y \in X$, if $d(x, y) < \delta$, then $e(f(x), f(y)) < \varepsilon$.

Then f is (d, e)-**uniformly continuous**.

Remarks

1. In particular, Theorem 3.3 tells us that every function in $C[0, 1]$ is uniformly continuous on $[0, 1]$ for the Euclidean metrics.
2. Notice that this definition does *not* require that the underlying metric space is sequentially compact.
3. The fact that a real-valued continuous function defined on a bounded closed interval (with the Euclidean metrics) is uniformly continuous can be used to prove that it is (Riemann) integrable.

This is proved in unit AB3 of M208.

We finish this subsection by recording a useful lemma that shows that uniformly continuous functions are always continuous.

Lemma 3.6

Let (X, d) and (Y, e) be metric spaces and suppose that $f\colon X \to Y$. If f is (d, e)-uniformly continuous on X, then f is (d, e)-continuous on X.

Exercise 3.3

Prove Lemma 3.6.

Exercise 3.4

Give an example of a continuous function $f\colon \mathbb{R} \to \mathbb{R}$ that is not uniformly continuous (for the Euclidean metrics on both its domain and codomain).

3.2 Sequentially compact subsets of $C[0, 1]$

In Worked Exercise 2.4 and Exercise 2.5, we investigated sets of functions in $C[0, 1]$ that turned out not to be sequentially compact. In this subsection we state, and show how to use, a result that gives us sufficient conditions to guarantee that a subset of $C[0, 1]$ *is* sequentially compact.

Before we state our main result for this section, we need a couple of definitions, both of which apply to a *set* of functions.

Our first definition gives a notion of boundedness for collections of functions.

Definition 3.7 *Pointwise bounded*

Let A be a set of functions from X to $\mathbb{R}$. We say that the set A is **pointwise bounded** if for each $x \in X$, there is $M_x \geq 0$ such that for each $f \in A$, $|f(x)| \leq M_x$.

So each $f \in A$ has domain X and codomain $\mathbb{R}$.

Remarks

1. The constant M_x may depend on the point $x \in X$ chosen but cannot depend on the choice of function from A.
2. This definition may remind you a little of the definition of pointwise convergence that we gave in Book D, Chapter 15 inasmuch as it does not require there to be a metric defined for the domain of the functions and so, in particular, there are no assumptions about whether the functions are continuous.
3. If the set A consists of just a single function, $A = \{f\}$ say, then this definition is automatically satisfied – a set consisting of one function is *always* pointwise bounded. In particular, saying the set $A = \{f\}$ is pointwise bounded is *not* the same as saying that the function f is bounded.

Exercise 3.5

Let $g\colon \mathbb{R} \to \mathbb{R}$ be a function and let A_g be the set consisting of all functions $f\colon \mathbb{R} \to \mathbb{R}$ such that $|f(x)| \leq |g(x)|$ for each $x \in \mathbb{R}$. Show that A is pointwise bounded.

Exercise 3.6

Let $A = \{f\colon [0,1] \to \mathbb{R} : f \text{ is constant}\}$.

Show that A is not pointwise bounded.

Exercise 3.7

Let (X, d) be a metric space, let $L \geq 0$, and fix $a \in X$ and $b \in \mathbb{R}$. Define

$$A = \{f\colon X \to \mathbb{R} : f(a) = b \text{ and } f \text{ is } (d, d^{(1)})\text{-Lipschitz with Lipschitz constant } L\}.$$

Lipschitz functions are defined in Book D, Chapter 15. A function $f\colon X \to \mathbb{R}$ is $(d, d^{(1)})$-Lipschitz with Lipschitz constant L if $|f(x) - f(y)| \leq Ld(x, y)$ for each $x, y \in X$.

Show that A is a pointwise bounded set.

Exercise 3.8

Let (X, d) be a metric space and let $L \geq 0$. Define

$$A = \{f\colon X \to \mathbb{R} : f \text{ is } (d, d^{(1)})\text{-Lipschitz with Lipschitz constant } L\}.$$

Determine whether A is a pointwise bounded set.

Exercise 3.9

Show that if A is a sequentially compact subset of $(C[0,1], d_{\max})$, then it is pointwise bounded.

Our second definition extends the notion of uniform continuity to a collection of functions.

Definition 3.8 *Equicontinuity*

Let (X, d) and (Y, e) be metric spaces and let A be a set of functions from X to Y.

The set A is a (d, e)**-equicontinuous set of functions** if for each $\varepsilon > 0$, there is a $\delta > 0$ such that for each $f \in A$ and every $x, y \in X$, $e(f(x), f(y)) < \varepsilon$ whenever $d(x, y) < \delta$.

Remarks

1. Although we give the general definition here, we shall only use it for collections of continuous functions from $[0,1]$ to $\mathbb{R}$.
2. We usually suppress mention of the metrics when it is clear from the context and just say that A is an equicontinuous set (or family) of functions.
3. If you compare this definition with that of the definition of uniform continuity (for a single function), Definition 3.5, you will see that the main difference is that the same δ has to work for all functions f in the set A. In particular, if A consists of just a single function $A = \{f\}$, then saying that A is equicontinuous is the same as saying that f is uniformly continuous.

Worked Exercise 3.9

Let $L \geq 0$ be given and let

$$A = \{f \colon [0,1] \to \mathbb{R} : f \text{ is } (d^{(1)}_{[0,1]}, d^{(1)})\text{-Lipschitz with Lipschitz constant } L\}.$$

Show that A is an equicontinuous set of functions (for the Euclidean metrics on $[0,1]$ and $\mathbb{R}$).

Solution

Since each $f \in A$ is Lipschitz with Lipschitz constant L, we have for each $x, y \in [0,1]$,

$$|f(x) - f(y)| \leq L|x - y|.$$

The '$L+1$' is to cover the case $L=0$.

Hence if $|x-y|<\delta$, then $|f(x)-f(y)|\leq L\delta$. Thus if $\varepsilon>0$ is given, then on choosing $\delta=\varepsilon/(L+1)$, we find that if $|x-y|<\delta$, then for any $f\in A$, $|f(x)-f(y)|<\varepsilon$. Hence A is an equicontinuous set of functions.

Exercise 3.10

Let (X,d) and (Y,e) be metric spaces and let A be a (d,e)-equicontinuous set of functions from X to Y. Show that if $B\subseteq A$, then B is also a (d,e)-equicontinuous set of functions.

Our next lemma shows that functions in an equicontinuous set are continuous.

Lemma 3.10

Let (X,d) and (Y,e) be metric spaces and let A be a (d,e)-equicontinuous set of functions from X to Y. Then each $f\in A$ is (d,e)-uniformly continuous and hence (d,e)-continuous on X.

Exercise 3.11

Prove Lemma 3.10.

We now identify a class of subsets of $C[0,1]$ that are equicontinuous (for the Euclidean metrics on both domain and codomain).

Proposition 3.11

If A is a $d_{\max}$-sequentially compact subset of $C[0,1]$, then A is equicontinuous.

Proof We give a proof by contradiction: suppose that A is a $d_{\max}$-sequentially compact subset of $C[0,1]$ that is not equicontinuous.

Then there is $\varepsilon>0$ such that for each $\delta>0$, there is $f\in A$ and $x,y\in[0,1]$ such that $|f(x)-f(y)|\geq\varepsilon$ and $|x-y|<\delta$.

In particular, taking $\delta=\frac{1}{n}$, we can find $f_n\in A$ and $x_n,y_n\in[0,1]$ such that $|f_n(x_n)-f_n(y_n)|\geq\varepsilon$ and $|x_n-y_n|\leq\frac{1}{n}$. Thus we have a sequence of functions (f_n) and two sequences of points (x_n) and (y_n). Since A is sequentially compact, (f_n) has a $d_{\max}$-convergent subsequence, (f_{n_k}), say, converging to a function $f\in A$. In particular, f is continuous.

Consider now the corresponding sequence of points (x_{n_k}), which lies within the interval $[0,1]$. Since, by Theorem 1.2, $[0,1]$ is $d^{(1)}$-sequentially compact, we can find a subsequence of (x_{n_k}), $(x_{n_{k_i}})_{i=1}^{\infty}$, say, that converges to $x \in [0,1]$. Moreover, since $|x_n - y_n| < \frac{1}{n}$ for each $n \in \mathbb{N}$, we conclude that the sequence $(y_{n_{k_i}})_{i=1}^{\infty}$ also converges to x. Now choose $I \in \mathbb{N}$ so that $d_{\max}(f_{n_k}, f) < \frac{\varepsilon}{4}$ for $k > I$. Then since $k_i \geq i$, we have that for $i > I$,

$$\begin{aligned}
\varepsilon &\leq |f_{n_{k_i}}(x_{n_{k_i}}) - f_{n_{k_i}}(y_{n_{k_i}})| \\
&= |f_{n_{k_i}}(x_{n_{k_i}}) - f(x_{n_{k_i}}) + f(x_{n_{k_i}}) - f(x) + f(x) - f(y_{n_{k_i}}) + f(y_{n_{k_i}}) - f_{n_{k_i}}(y_{n_{k_i}})| \\
&\leq |f_{n_{k_i}}(x_{n_{k_i}}) - f(x_{n_{k_i}})| + |f(x_{n_{k_i}}) - f(x)| + |f(x) - f(y_{n_{k_i}})| \\
&\qquad + |f(y_{n_{k_i}}) - f_{n_{k_i}}(y_{n_{k_i}})| \\
&\leq d_{\max}(f, f_{n_{k_i}}) + |f(x_{n_{k_i}}) - f(x)| + |f(x) - f(y_{n_{k_i}})| + d_{\max}(f, f_{n_{k_i}}) \\
&< \tfrac{\varepsilon}{2} + |f(x_{n_{k_i}}) - f(x)| + |f(x) - f(y_{n_{k_i}})|.
\end{aligned}$$

Since f is continuous and $x_{n_{k_i}} \to x$, we can choose $J > I$ so that for $i > J$, $|f(x_{n_{k_i}}) - f(x)| < \frac{\varepsilon}{4}$. Hence for $i > J$, we have

$$\varepsilon < \frac{3\varepsilon}{4} + |f(x) - f(y_{n_{k_i}})|,$$

which rearranges to give $|f(x) - f(y_{n_{k_i}})| > \frac{\varepsilon}{4}$ for $i > J$. But $y_{n_{k_i}} \to x$ as $i \to \infty$, and so we conclude that f is not continuous at x, which is a contradiction. ■

Summarising our results for sequentially compact subsets of $C[0,1]$ (with the $d_{\max}$ metric), we have:

1. sequentially compact subsets of $C[0,1]$ are $d_{\max}$-closed (by Theorem 2.5)
2. sequentially compact subsets of $C[0,1]$ are pointwise bounded (by Exercise 3.9)
3. sequentially compact subsets of $C[0,1]$ are equicontinuous (by Proposition 3.11).

The following theorem, which is the main result of this section, shows that the converse also holds: these three properties are sufficient to identify sequentially compact subsets of $(C[0,1], d_{\max})$.

Theorem 3.12 *The Arzelà–Ascoli Theorem*

Let A be a subset of $C[0,1]$. Suppose that:

1. A is $d_{\max}$-closed
2. A is a pointwise bounded set of functions
3. A is an equicontinuous set of functions.

Then A is $d_{\max}$-sequentially compact.

This theorem is named after the Italian mathematicians Cesare Arzelà (1847–1912) and Giulio Ascoli (1843–1896), who both proved versions of it. Variants of this theorem play an important role in mathematical analysis.

The proof is long and difficult and we defer it to Subsection 3.3; we spend the remainder of this subsection seeing how we can use the theorem.

Worked Exercise 3.13

Let

$$A = \{f \in C[0,1] : f(0) = 0 \text{ and } |f(x) - f(y)| \leq |x - y| \text{ for each } x, y \in [0,1]\}.$$

Define $I : C[0,1] \to \mathbb{R}$ by

$$I(f) = \int_0^1 f(x)\,dx.$$

Show that I attains a minimum on A.

Solution

For example, A contains the function $f : [0,1] \to \mathbb{R}$ given by $f(x) = 0$ for $x \in [0,1]$.

We first note that the set A is non-empty. In Book D, Chapter 15, Worked Exercise 5.6, we showed that I is a $(d_{\max}, d^{(1)})$-continuous function. If we can prove that A is a sequentially compact set, then the General Extreme Value Theorem (Theorem 2.13) tells us that I is bounded on A and attains its bounds. In particular, there is an $f_0 \in A$ such that

$$I(f_0) \leq I(f) \quad \text{for all } f \in A,$$

and so I attains its minimum on A.

In fact, f_0 is the function given by $f_0(x) = -x$, but we are not asked to show this.

Thus it only remains for us to show that A is sequentially compact. To do this, we use Theorem 3.12 and verify that:

1. A is $d_{\max}$-closed
2. A is a pointwise bounded family of functions
3. A is an equicontinuous family of functions.

A is closed

Let (f_n) be a sequence of functions in A that $d_{\max}$-converges to $f \in C[0,1]$. To show that A is closed, we must show that $f \in A$.

Since $f_n(0) = 0$ for each $n \in \mathbb{N}$, we have

$$|f(0)| = |f(0) - f_n(0)| \leq d_{\max}(f_n, f).$$

But $d_{\max}(f_n, f)$ is a null sequence. Hence $f(0) = 0$.

For each $x, y \in [0,1]$, we can use the Triangle Inequality (for the modulus) twice and the fact that $f_n \in A$ to estimate that for each $n \in \mathbb{N}$,

$$\begin{aligned} |f(x) - f(y)| &\leq |f(x) - f_n(x)| + |f_n(x) - f_n(y)| + |f_n(y) - f(y)| \\ &\leq |f_n(x) - f_n(y)| + 2d_{\max}(f_n, f) \\ &\leq |x - y| + 2d_{\max}(f_n, f). \end{aligned}$$

But $n \in \mathbb{N}$ is arbitrary and $d_{\max}(f_n, f) \to 0$ as $n \to \infty$, so

$$|f(x) - f(y)| \leq |x - y| \quad \text{for each } x, y \in [0,1].$$

Hence f is 1-Lipschitz and $f(0) = 0$, and so $f \in A$, implying that A is closed.

A is a pointwise bounded family of functions

This follows immediately from the result of Exercise 3.7 (with (X, d) equal to $[0, 1]$ with the Euclidean metric, $a = 0$ and $L = 1$).

A is an equicontinuous family of functions

Since each function in A is $(d^{(1)}_{[0,1]}, d^{(1)})$-Lipschitz with Lipschitz constant 1, we can use the result of Worked Exercise 3.9 together with that of Exercise 3.10 to conclude that A is an equicontinuous set of functions.

Since the hypotheses of Theorem 3.12 hold, we conclude that A is sequentially compact, as required.

Exercise 3.12

Let

$$A = \{f \in C[0,1] : f(0) = 0 \text{ and } |f(x) - f(y)| \leq \sqrt{|x-y|} \text{ for all } x, y \in [0,1]\}.$$

Define $F\colon C[0,1] \to \mathbb{R}$ by

$$F(f) = \int_0^1 x f(x)\, dx.$$

(a) Prove that A is $d_{\max}$-sequentially compact.

(b) Prove that F attains a minimum on A.

3.3 Proof of Theorem 3.12

This subsection is not assessed.

In this subsection, we prove Theorem 3.12. The proof is long and involved and you should not be surprised if you find it difficult – your goal should be to understand the outline of the argument and you need not be too concerned with trying to follow every single detail.

For the remainder of this subsection, we suppose that A is a subset of $C[0, 1]$ that has the following properties:

1. A is $d_{\max}$-closed
2. A is pointwise bounded – that is, for each $x \in [0, 1]$, there is an $M_x \geq 0$ such that $|f(x)| \leq M_x$ for each $f \in A$
3. A is equicontinuous – that is, for each $\varepsilon > 0$, there is a $\delta > 0$ such that for each $f \in A$ and each $x, y \in [0, 1]$, $|f(x) - f(y)| < \varepsilon$ whenever $|x - y| < \delta$.

Our aim is to show that A is $d_{\max}$-sequentially compact – that is, each sequence of functions in A has a $d_{\max}$-convergent subsequence whose limit is also in A.

So let (f_n) be a sequence of functions in A. Our objective is to 'build' a continuous function $f \in C[0,1]$ to which a subsequence of (f_n) $d_{\max}$-converges. We do this in two stages:

1. we find a candidate function f to which a subsequence of (f_n) may converge
2. we show that there is a subsequence of (f_n) that converges to this function f, and conclude that, since A is closed, $f \in A$.

Before embarking on the proof proper, we outline the main ideas.

Sketch of argument

- We start by observing that, since A is pointwise bounded, for each $a \in [0,1]$, there is a real number $M_a \geq 0$ for which $|f_n(a)| \leq M_a$ for each $n \in \mathbb{N}$. In particular, the sequence of real numbers $(f_n(a))$ lies in the $d^{(1)}$-sequentially compact interval $[-M_a, M_a]$ (Corollary 2.2) and so has a convergent subsequence, $(f_{n_k}(a))_{k=1}^{\infty}$ say, that converges to a real number in $[-M_a, M_a]$. Notice, however, that the choice of subsequence depends on the choice of a – different points may require different subsequences, and this is where the difficulty lies.
- But it is possible to find a convergent subsequence of (f_n) that works for two points a and b, say, simultaneously. First, given a, we can find an increasing sequence $1 \leq n_1 < n_2 < \cdots$ of integers so that $(f_{n_k}(a))_{k=1}^{\infty}$ is a convergent sequence converging to a real number, $f(a)$, say. By now considering $(f_{n_k})_{k=1}^{\infty}$ as a sequence of functions in its own right, we can find a subsequence $(n_{k_i})_{i=1}^{\infty}$ of (n_k) so that $(f_{n_{k_i}}(b))$ is also a convergent sequence of real numbers, converging to a real number $f(b)$, say. But $(f_{n_{k_i}}(a))$, being a subsequence of the convergent sequence $(f_{n_k}(a))_{k=1}^{\infty}$, is convergent to $f(a)$, and so we have found a subsequence of the functions (f_n) that converges for both a and b.
- By careful use of induction, we should be able to extend this argument to cope with a countable set of points $\{a_i : i \in \mathbb{N}\}$ (by taking a sequence of subsequences of previously found subsequences).
- Unfortunately, as we saw in Book D, Chapter 16, Theorem 6.7, the interval $[0,1]$ is not countable and so we cannot hope to deal with every point in $[0,1]$ in this way. However, if the function f that we are constructing is to be continuous on $[0,1]$, then we need only find its values on a set that is dense in the interval $[0,1]$, since then the value of f at an arbitrary point a in $[0,1]$ can be found by calculating $\lim_{i\to\infty} f(a_i)$, where (a_i) is a sequence of points lying in the dense set that converges to the point a. This provides us with the key idea needed for proving the theorem: use a *countable* dense subset of $[0,1]$ together with the 'taking subsequences' idea described above to define f on a subset of $[0,1]$, and then show that the f thus defined can be extended to the whole of $[0,1]$ in such a way to have the properties we want – f is continuous, f is a member of A and f is a limit of a subsequence of (f_n).

Book D, Chapter 16, Definition 6.1.

Finding a candidate limit function f

We first identify an appropriate countable dense subset of $[0,1]$ on which we hope to define f via repeated use of subsequences.

Let (a_i) be a sequence of real numbers whose set of values $\{a_i : i \in \mathbb{N}\}$ is dense in $[0,1]$ – that is,

$$\mathrm{Cl}_{(\mathbb{R},d^{(1)})}(\{a_i : i \in \mathbb{N}\}) = [0,1].$$

An example of such a sequence is provided by

$$\tfrac{1}{2}, \tfrac{1}{3}, \tfrac{2}{3}, \tfrac{1}{4}, \tfrac{2}{4}, \tfrac{3}{4}, \ldots, \tfrac{1}{n}, \tfrac{2}{n}, \ldots, \tfrac{n-1}{n}, \tfrac{1}{n+1}, \ldots,$$

This sequence takes every rational value in $(0,1)$ a (countably) infinite number of times.

but there are many other possibilities.

In order to start defining our function f, we try to find a strictly increasing sequence $(n_k)_{k\in\mathbb{N}}$ such that $(f_{n_k}(a_i))_{k=1}^{\infty}$ is a $d^{(1)}$-convergent sequence for each $i \in \mathbb{N}$. This will then allow us to define f on the set of values of the sequence (a_i) by setting $f(a_i) = \lim_{k\to\infty} f_{n_k}(a_i)$ for each $i \in A$.

We construct such a sequence by induction.

We start by choosing a strictly increasing sequence of natural numbers $(n_{1,k})_{k\in\mathbb{N}}$ for which $(f_{n_{1,k}}(a_1))_{k\in\mathbb{N}}$ is a $d^{(1)}$-convergent sequence.

We have shown earlier that we can do this.

Now suppose that for some $j \in \mathbb{N}$, we have a strictly increasing sequence $(n_{j,k})_{k\in\mathbb{N}}$, such that $(f_{n_{j,k}}(a_i))_{k\in\mathbb{N}}$ is a convergent sequence of real numbers for $i = 1, 2, \ldots, j$. Then, since $[-M_{a_{j+1}}, M_{a_{j+1}}]$ is sequentially compact, we can find a subsequence $(n_{j+1,k})_{k\in\mathbb{N}}$ of $(n_{j,k})_{k\in\mathbb{N}}$ such that $(f_{n_{j+1,k}}(a_{j+1}))_{k\in\mathbb{N}}$ is also $d^{(1)}$-convergent.

This subsequence is also strictly increasing.

Proceeding by induction, we find a sequence of strictly increasing sequences $((n_{j,k})_{k\in\mathbb{N}})_{j\in\mathbb{N}}$ for which:

1. $(n_{j+1,k})_{k\in\mathbb{N}}$ is a subsequence of $(n_{j,k})_{k\in\mathbb{N}}$, for each $j \in \mathbb{N}$
2. $n_{k,k} < n_{k+1,k+1}$ for each $k \in \mathbb{N}$, since $n_{k,k} \le n_{k+1,k} < n_{k+1,k+1}$
3. for each $j \in \mathbb{N}$, $(f_{n_{j,k}}(a_i))_{k\in\mathbb{N}}$ is convergent for $i = 1, 2, \ldots, j$.

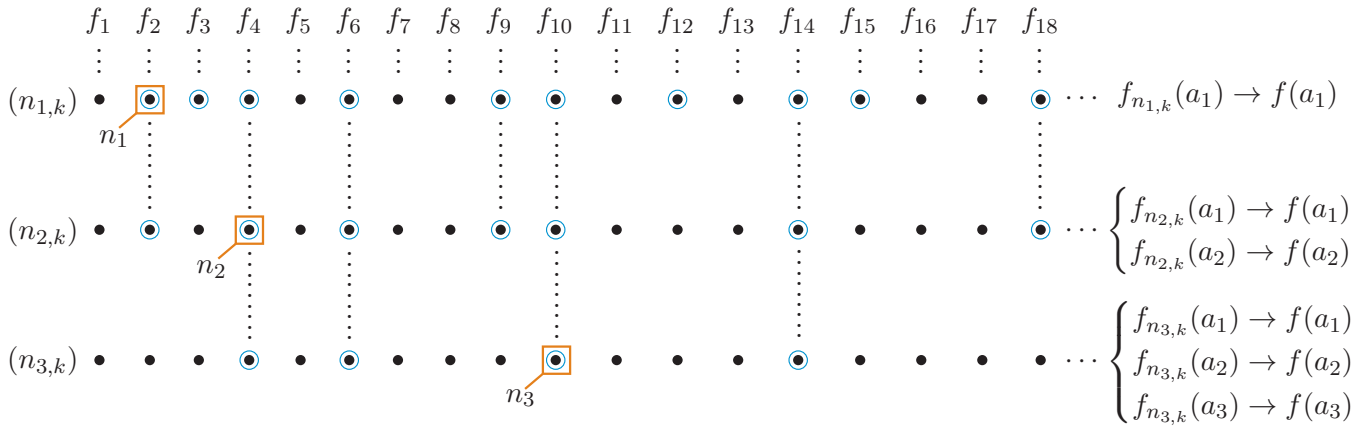

Figure 3.1 This figure shows how we construct the diagonal sequence (n_k), ensuring that the sequence $(f_{n_k}(a_i))_{k=1}^{\infty}$ is convergent for each i

In Book D, Chapter 16, we used a diagonalisation argument to show that $(0,1)$ is uncountable.

We now consider the diagonal sequence $(n_k)_{k=1}^{\infty}$, where $n_k = n_{k,k}$. Then, for each $i \in \mathbb{N}$, $(f_{n_k}(a_i))_{k\in\mathbb{N}}$ is a $d^{(1)}$-convergent sequence, since $(n_k)_{k=i}^{\infty}$ is a subsequence of $(n_{i,k})_{k\in\mathbb{N}}$ and $(f_{n_{i,k}}(a_i))_{k=1}^{\infty}$ is convergent (see Figure 3.1).

We now start to define the function f that will turn out to be the $d_{\max}$-limit of the sequence (f_{n_k}). We begin by defining its value for each a_i in our dense sequence (a_i) by setting

$$f(a_i) = \lim_{k\to\infty} f_{n_k}(a_i), \text{ for each } i \in \mathbb{N}.$$

This function is well defined because if $a_i = a_{i+j}$ for some $j \in \mathbb{N}$, then

$$f(a_i) = \lim_{k\to\infty} f_{n_k}(a_i) = \lim_{k\to\infty} f_{n_k}(a_{i+j}) = f(a_{i+j}).$$

Thus we have defined a function $f\colon \{a_i : i \in \mathbb{N}\} \to \mathbb{R}$.

It remains to extend the definition of f to arbitrary points in $[0,1]$, and this is where we use the fact that $\{a_i : i \in \mathbb{N}\}$ is dense in $[0,1]$. For each $a \in [0,1]$, there is a subsequence $(a_{i_p})_{p=1}^{\infty}$ of $(a_i)_{i=1}^{\infty}$ for which $a_{i_p} \to a$ as $p \to \infty$. We would like to use this subsequence to define $f(a)$ by setting

$$f(a) = \lim_{p\to\infty} f(a_{i_p}).$$

Unfortunately there are two potential problems:

(i) given a subsequence (a_{i_p}) of (a_i) that converges to a, it is not obvious that $(f(a_{i_p}))$ is a convergent sequence

(ii) even if $(f(a_{i_p}))$ is convergent, there may be another choice of subsequence $(a_{j_p})_{p=1}^{\infty}$ of (a_i) that converges to a for which $(f(a_{j_p}))$ does not converge to $\lim_{p\to\infty} f(a_{i_p})$ – this would give us a different inconsistent choice for the definition of $f(a)$.

Fortunately, the fact that A is equicontinuous enables us to show that neither problem can arise.

We first show that

Claim 1

If $(a_{i_p})_{p\in\mathbb{N}}$ and $(a_{j_p})_{p\in\mathbb{N}}$ are subsequences of (a_i) that both converge to a, then $|f(a_{i_p}) - f(a_{j_p})| \to 0$ as $p \to \infty$.

Proof Let $\varepsilon > 0$ be given: we show that there is $P \in \mathbb{N}$ such that for $p > P$, $|f(a_{i_p}) - f(a_{j_p})| < \varepsilon$, implying the claim.

Since A is equicontinuous, we can choose a $\delta > 0$ so that for each $g \in A$ and each $x, y \in [0,1]$,

$$|g(x) - g(y)| < \tfrac{1}{4}\varepsilon \quad \text{whenever } x, y \in [0,1] \text{ satisfy } |x - y| < \delta. \tag{1}$$

Using the Triangle Inequality (for modulus) three times implies that for each $n \in \mathbb{N}$,

$$
\begin{aligned}
&|f(a_{i_p}) - f(a_{j_p})| \\
&= |f(a_{i_p}) - f_n(a_{i_p}) + f_n(a_{i_p}) - f_n(a) + f_n(a) - f_n(a_{j_p}) + f_n(a_{j_p}) - f(a_{j_p})| \\
&\leq |f(a_{i_p}) - f_n(a_{i_p})| + |f_n(a_{i_p}) - f_n(a)| \\
&\qquad + |f_n(a) - f_n(a_{j_p})| + |f_n(a_{j_p}) - f(a_{j_p})|.
\end{aligned}
$$

Since (a_{i_p}) and (a_{j_p}) both converge to a as $p \to \infty$, there is $P \in \mathbb{N}$ such that for $p > P$, both $|a_{i_p} - a| < \delta$ and $|a_{j_p} - a| < \delta$. For $p > P$, we then have by equation (1), for *any* choice of $n \in \mathbb{N}$,

$$|f_n(a_{i_p}) - f_n(a)| < \tfrac{1}{4}\varepsilon \quad \text{and} \quad |f_n(a) - f_n(a_{j_p})| < \tfrac{1}{4}\varepsilon.$$

For each $p \in \mathbb{N}$, $f_{n_k}(a_{i_p}) \to f(a_{i_p})$ and $f_{n_k}(a_{j_p}) \to f(a_{j_p})$ as $k \to \infty$, and so for each $p > P$ there is an $K \in \mathbb{N}$ for which

$$|f(a_{i_p}) - f_{n_K}(a_{i_p})| < \tfrac{1}{4}\varepsilon \quad \text{and} \quad |f_{n_K}(a_{j_p}) - f(a_{j_p})| < \tfrac{1}{4}\varepsilon.$$

Hence, fixing some $p > P$ and choosing a $K \in \mathbb{N}$ so that these inequalities hold, we deduce that

$$|f(a_{i_p}) - f(a_{j_p})| \leq \tfrac{1}{4}\varepsilon + \tfrac{1}{4}\varepsilon + \tfrac{1}{4}\varepsilon + \tfrac{1}{4}\varepsilon = \varepsilon.$$

Since $\varepsilon > 0$ is arbitrary, we conclude that

$$|f(a_{i_p}) - f(a_{j_p})| \to 0 \text{ as } p \to \infty, \text{ as claimed.} \tag{2}$$

■

We now show that

Claim 2

If $a_{i_p} \to a$ as $p \to \infty$, then $(f(a_{i_p}))_{p\in\mathbb{N}}$ is $d^{(1)}$-convergent.

Proof Suppose that $a_{i_p} \to a$ as $p \to \infty$. We first show that $(f(a_{i_p}))$ is a bounded sequence: by using the Triangle Inequality (for modulus) twice, we find that for each $p, N \in \mathbb{N}$ we have

$$
\begin{aligned}
|f(a_{i_p})| &= |f(a_{i_p}) - f_n(a_{i_p}) + f_n(a_{i_p}) - f_n(a) + f_n(a)| \\
&\leq |f(a_{i_p}) - f_n(a_{i_p})| + |f_n(a_{i_p}) - f_n(a)| + |f_n(a)|.
\end{aligned}
$$

But A is pointwise bounded, and so $(f_n(a))$ is a bounded sequence and there is $M_a > 0$ for which $|f_n(a)| \leq M_a$ for each $n \in \mathbb{N}$. Moreover, A is equicontinuous, and so we can find $\delta > 0$ for which if $x, y \in [0,1]$ with $|x - a| < \delta$, then $|f_n(x) - f_n(y)| < 1$ for each $n \in \mathbb{N}$. Fix $P \in \mathbb{N}$ for which $|a_{i_p} - a| < \delta$ for $p > P$. Then for each $p > P$, we can find n_k so that $|f(a_{i_p}) - f_{n_k}(a_{i_p})| < 1$. Hence for such a $p > P$ and n_k,

$$|f(a_{i_p})| \leq 1 + 1 + M_a = 2 + M_a.$$

Thus for each $p \in \mathbb{N}$,

$$|f(a_{i_p})| \leq \max\{|f(a_{i_p})| : p = 1, 2, \ldots, P\} + 2 + M_a < \infty,$$

and so $(f(a_{i_p}))_{p=1}^{\infty}$ is a bounded sequence of real numbers. Thus by the Monotone Convergence Theorem we can find a convergent subsequence

$(f(a_{i_{p_j}}))_{j=1}^{\infty}$ with limit L, say. It now follows from the Triangle Inequality and equation (2) that

$$|f(a_{i_p}) - L| \leq |f(a_{i_p}) - f(a_{i_{p_j}})| + |f(a_{i_{p_j}}) - L|.$$

The right-hand side of this inequality can be made arbitrarily small by choosing p and j sufficiently large, and so $\lim_{p\to\infty} f(a_{i_p})$ exists (and is L). ■

Claim 1 tells us that if (a_{i_p}) and (a_{j_p}) both converge to a and $(f(a_{i_p}))$ is convergent, then $f((a_{j_p}))$ is also convergent and converges to the same value as $(f(a_{i_p}))$. Hence we can define the value of f at a to be $\lim_{p\to\infty} f(a_{i_p})$, where a_{i_p} is any subsequence of (a_i) that converges to a. This completes our definition of the function $f\colon [0,1] \to \mathbb{R}$.

A subsequence of (f_n) converges to f

It only remains to show that $f \in A$: we need to show that f is continuous (and so in $C[0,1]$) and there is a sequence of functions (f_n) from A that converges to f for $d_{\max}$. We do this by using the notion of *uniform convergence* and show that $f_{n_k} \to f$ uniformly as $k \to \infty$ (where n_k is the strictly increasing sequence we found previously). It then follows from Book D, Chapter 15, Theorem 5.8 that $f \in C[0,1]$, with $f_{n_k} \overset{d_{\max}}{\to} f$ as $k \to \infty$ and so $f \in A$, since A is $d_{\max}$-closed.

Recall from Book D, Chapter 15, Definition 4.4 that a sequence (f_k) of functions $f_k\colon [0,1] \to \mathbb{R}$ **converges uniformly** on $[0,1]$ to the function $f\colon [0,1] \to \mathbb{R}$ if the function $f_k - f$ is bounded for each $k \in \mathbb{N}$, and the sequence (M_k) defined by $M_k = \sup\{|f_k(x) - f(x)| : x \in [0,1]\}$ is a (real) null sequence.

We split the proof that $f_{n_k} \to f$ uniformly into three steps:

(1) we show that (f_{n_k}) converges to f pointwise, that is, for each $a \in [0,1]$, $f_{n_k}(a) \to f(a)$ as $k \to \infty$

(2) we show that f is uniformly continuous and so, by Lemma 3.6, f is continuous

(3) we use steps (1) and (2) to show that $f_{n_k} \to f$ uniformly as $k \to \infty$.

Step (1)

$f_{n_k}(a) \to f(a)$ as $k \to \infty$ for each $a \in [0,1]$.

Let $a \in [0,1]$, and let $\varepsilon > 0$ be given. We must show that there is a $K \in \mathbb{N}$ such that $|f(a) - f_{n_k}(a)| < \varepsilon$ for each $k > K$.

Since A is equicontinuous, we can certainly choose $\delta > 0$ so that for each $x, y \in [0,1]$ if $|x - y| < \delta$, then $|f_n(x) - f_n(y)| < \frac{1}{3}\varepsilon$ for every $n \in \mathbb{N}$. In particular, whenever $x \in [0,1]$ with $|x - a| < \delta$, then $|f_n(x) - f_n(a)| < \frac{1}{3}\varepsilon$ for any $n \in \mathbb{N}$.

Since $(a_i)_{i\in\mathbb{N}}$ is dense in $[0,1]$, there is a subsequence $(a_{i_p})_{p\in\mathbb{N}}$ for which $a_{i_p} \to a$ as $p \to \infty$. And so, by Claim 1, $f(a_{i_p}) \to f(a)$ as $p \to \infty$.

Hence we can find a $P \in \mathbb{N}$ such that for $p > P$, $|a_{i_p} - a| < \delta$ and $|f(a_{i_p}) - f(a)| < \frac{1}{3}\varepsilon$.

On fixing such a $p > P$, using the Triangle Inequality (for modulus) twice gives

$$\begin{aligned}|f(a) - f_{n_k}(a)| &\leq |f(a) - f(a_{i_p})| + |f(a_{i_p}) - f_{n_k}(a_{i_p})| + |f_{n_k}(a_{i_p}) - f_{n_k}(a)| \\ &< \tfrac{1}{3}\varepsilon + |f(a_{i_p}) - f_{n_k}(a_{i_p})| + \tfrac{1}{3}\varepsilon.\end{aligned}$$

But $f_{n_k}(a_{i_p}) \to f(a_{i_p})$ as $k \to \infty$, and so there is $K \in \mathbb{N}$ (that depends on p) so that $|f(a_{i_p}) - f_{n_k}(a_{i_p})| < \frac{1}{3}\varepsilon$ for each $k > K$. Hence for $k > K$,

$$|f(a) - f_{n_k}(a)| < \varepsilon.$$

Since ε is arbitrary, we conclude that $f_{n_k}(a) \to f(a)$ as $k \to \infty$.

Step (2)

f is uniformly continuous (and so is continuous).

Let $\varepsilon > 0$ be given.

Since A is an equicontinuous set, we can find $\delta > 0$ so that, for each $a, b \in [0,1]$ with $|b - a| < \delta$, we have for every $n \in \mathbb{N}$, $|f_n(b) - f_n(a)| < \frac{1}{3}\varepsilon$.

For any $a, b \in [0,1]$, by using the Triangle Inequality (for modulus) twice we find that for each $k \in \mathbb{N}$,

$$\begin{aligned}|f(b) - f(a)| &\leq |f(b) - f_{n_k}(b)| + |f_{n_k}(b) - f_{n_k}(a)| + |f_{n_k}(a) - f(a)| \\ &< |f(b) - f_{n_k}(b)| + \tfrac{1}{3}\varepsilon + |f_{n_k}(a) - f(a)|.\end{aligned}$$

But, by the result of Step (1), $f_{n_k}(b) \to f(b)$ and $f_{n_k}(a) \to f(a)$ as $k \to \infty$. Hence there is a $K \in \mathbb{N}$ so that, for $k > K$, both $|f(b) - f_{n_k}(b)| < \frac{1}{3}\varepsilon$ and $|f_{n_k}(a) - f(a)| < \frac{1}{3}\varepsilon$. Thus on fixing such a $k > K$, we find

$$|f(b) - f(a)| < \tfrac{1}{3}\varepsilon + \tfrac{1}{3}\varepsilon + \tfrac{1}{3}\varepsilon = \varepsilon.$$

That is, f is uniformly continuous.

Step (3)

$f_{n_k} \to f$ uniformly as $k \to \infty$.

Since f and f_{n_k} are continuous functions, their difference is bounded for each k. It remains to show that (M_k) given by $M_k = \sup\{|f(a) - f_{n_k}(a)| : a \in [0,1]\}$ is a null sequence. To do this, it is enough to show that for each $\varepsilon > 0$, there is a $K \in \mathbb{N}$ so that $|f(a) - f_{n_k}(a)| < \varepsilon$, for each $k > K$ and each $a \in [0,1]$.

Let $\varepsilon > 0$ be given. Since A is equicontinuous and (by Step (2)) f is uniformly continuous, there is $\delta > 0$ so that if $a, b \in [0,1]$ with $|a - b| < \delta$, then both $|f(b) - f(a)| < \frac{1}{3}\varepsilon$ and for each $n \in \mathbb{N}$, $|f_n(b) - f_n(a)| < \frac{1}{3}\varepsilon$.

Let

$$B = \{j\delta : j = 0, 1, 2, \ldots\} \cap [0,1].$$

(So B is a finite set and each $a \in [0,1]$ is at most a distance δ away from a point in B.)

The result of Step (1) tells us that for each $b \in B$, we can find $K_b \in \mathbb{N}$ so that $|f(b) - f_{n_k}(b)| < \frac{1}{3}\varepsilon$ for $k > K_b$.

Let $K = \max\{K_b : b \in B\}$, then K exists and is finite, since B is finite.

If $k > K$, then for each $a \in A$, on choosing $b \in B$ so that $|b - a| < \delta$, we can use the Triangle Inequality (for modulus) twice, to estimate that

$$\begin{aligned} |f(a) - f_{n_k}(a)| &\leq |f(a) - f(b)| + |f(b) - f_{n_k}(b)| + |f_{n_k}(b) - f_{n_k}(a)| \\ &< \tfrac{1}{3}\varepsilon + \tfrac{1}{3}\varepsilon + \tfrac{1}{3}\varepsilon = \varepsilon. \end{aligned}$$

Hence $f_{n_k} \to f$ uniformly as $k \to \infty$ and we are done.

4 Compact metric spaces

So far we have studied *sequential* compactness. It is natural to wonder why the term 'compactness' is qualified by the word 'sequential' (as mathematicians are not known for using two words when one will suffice). You may have encountered unqualified use of the term 'compactness' in the *History Reader*, and you may have wondered whether the two terms are the same. In this section we define the notion of a *compact* metric space, and show that for a metric space, the two concepts are equivalent.

Compactness is defined in terms of the open sets of a space. We originally defined open sets as being the complements of closed sets but, as with many mathematical concepts, there are alternative characterisations, and the open ball characterisation, given in Book D, Chapter 16, Theorem 2.9 can be more helpful when thinking about compactness. We restate the result here for reference.

Let (X, d) be a metric space and let A be a subset of X. Then A is d-open if, and only if, for each $a \in A$ there is $r > 0$ so that $B_d(a, r) \subseteq A$.

In other words, a d-open set is a set that contains a d-open ball around every one of its points (Figure 4.1). This characterisation can be useful when proving directly that a set is d-open.

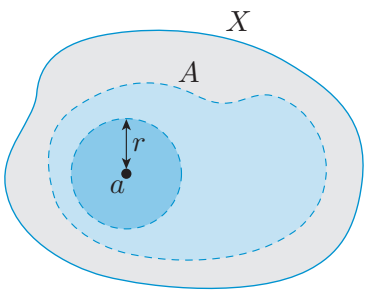

Figure 4.1

We showed that all subsets of a discrete metric space are open in Example 2.5 of Chapter 16, using the fact that an open set is the complement of a closed set.

Worked Exercise 4.1

Use the open set characterisation to prove that the set $\{p\}$, containing the single point $p \in [0,1]$, is d_0-open in $[0,1]$ when $[0,1]$ is considered as a metric space with the discrete metric.

Solution

We need to find an r such that $B_{d_0}(p, r) \subseteq \{p\}$. Take $r = \frac{1}{2}$. Then $B_{d_0}(p, r) = \{x \in [0, 1] \text{ such that } d_0(x, p) < \frac{1}{2}\} = \{p\}$ since $d_0(x, p)$ is equal to 1 for any $x \neq p$. Therefore $B_{d_0}(p, \frac{1}{2}) \subseteq \{p\}$. (In fact, any r with $0 < r \leq 1$ will suffice.)

Exercise 4.1

Use the open set characterisation to show that (a, b) is d-open in $\mathbb{R}$, where d is the metric induced by the usual Euclidean metric on $\mathbb{R}$.

Compactness, as is the case with sequential compactness, is a property of the whole space. Therefore we do not focus our attention on single open sets, but consider collections of open sets that 'cover' the whole space. By this we mean a collection of open sets, $\{U_i : i \in I\}$, which is such that each element of our space is contained within at least one of these open sets. Formally we define an *open cover* as follows.

Definition 4.2 *Open cover*

Let (X, d) be a metric space and let $A \subseteq X$.

A collection $\mathcal{S}$ of d-open subsets of X is an **open cover** of A if

$$A \subseteq \bigcup_{U \in \mathcal{S}} U;$$

that is, for each $a \in A$, there exists $U \in \mathcal{S}$ such that $a \in U$; see Figure 4.2.

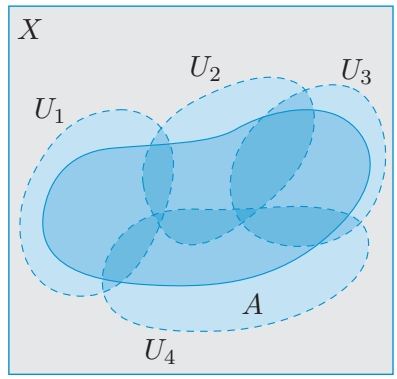

Figure 4.2

Remark

If we wish to emphasise the metric, we say d-open cover.

Let $A = [0, 1]$ with the usual metric, $d^{(1)}$, induced by the Euclidean metric on $\mathbb{R}$. There are many different open covers of A. We begin with the following example.

Worked Exercise 4.3

Show that

$$\mathcal{S}_1 = \{(x - \tfrac{1}{10}, x + \tfrac{1}{10}) : x \in [0, 1]\}$$

is a $d^{(1)}$-open cover of $[0, 1]$, where $d^{(1)}$ is the usual Euclidean metric on $\mathbb{R}$.

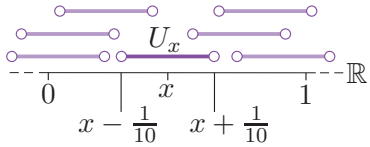

Figure 4.3

Solution

Let $U_x = (x - \frac{1}{10}, x + \frac{1}{10})$. If $x \in [0,1]$, then $x \in U_x$ and so

$$[0,1] \subseteq \bigcup_{x \in [0,1]} U_x = \bigcup_{U_x \in \mathcal{S}_1} U_x \quad \text{(Figure 4.3)}.$$

Also, each set U_x in $\mathcal{S}_1$ is d-open. So $\mathcal{S}_1$ is an open cover of $[0,1]$.

You may have noticed that many of the sets in $\mathcal{S}_1$ overlap, suggesting that $[0,1]$ can be covered by a smaller collection of sets. In fact, $[0,1]$ can be covered by just *finitely* many of the sets in $\mathcal{S}_1$!

Exercise 4.2

Find a finite subcollection of sets in $\mathcal{S}_1$ that covers $[0,1]$.

Now consider

$$\mathcal{S} = \{(0,q) : q \in \mathbb{Q} \cap (0,\infty)\}.$$

This is a collection of $d^{(1)}$-open sets with respect to the Euclidean metric, $d^{(1)}$ on $\mathbb{R}$, but it is not a cover of $[0,1]$, for no set in $\mathcal{S}$ contains the point 0. However,

$$\mathcal{S}_2 = \mathcal{S} \cup \{(-\tfrac{1}{100}, \tfrac{1}{100})\}$$

is an open cover of $[0,1]$.

Exercise 4.3

(a) Show that $\mathcal{S}_2$ is an open cover of $[0,1]$.

(b) Find a finite subcollection of sets in $\mathcal{S}_2$ that covers $[0,1]$.

In fact, it turns out that *any* $d^{(1)}$-open cover of $[0,1]$, contains a finite subcollection of sets that covers $[0,1]$. Before proving this result, we make the following definition.

Definition 4.4 *Finite subcover*

Let (X,d) be a metric space and let $A \subseteq X$. Let $\mathcal{S}$ be a d-open cover of A. A collection $\mathcal{R}$ of d-open subsets of X is a **finite (d-open) subcover** of A from $\mathcal{S}$ if:

(a) $\mathcal{R} \subseteq \mathcal{S}$

(b) $A \subseteq \bigcup_{U \in \mathcal{R}} U$

(c) $\mathcal{R}$ is a finite collection of sets.

Theorem 4.5

Every $d^{(1)}$-open cover of the interval $[0,1]$ has a finite subcover.

Proof Let $\mathcal{S}$ be any $d^{(1)}$-open cover of $[0,1]$. In other words, $\mathcal{S}$ is any collection of $d^{(1)}$-open sets of $\mathbb{R}$ which covers $[0,1]$. We must show that $\mathcal{S}$ contains a finite subcover of $[0,1]$.

Let A be the set of points $x \in [0,1]$ for which the interval $[0,x]$ is covered by a finite subcollection of sets from $\mathcal{S}$.

Considering A is not an obvious thing to do, but it lies at the heart of the proof.

Certainly $0 \in A$, since 0 belongs to at least one set from $\mathcal{S}$ (otherwise $\mathcal{S}$ would not be an open cover of $[0,1]$) and so that set (alone) is a finite subcover of 0. Hence A is not empty.

Observe that if $x \in A$ and $0 \leq y \leq x$, then $y \in A$ since, if $\{U_1, U_2, \ldots, U_n\}$ is a finite subset of $\mathcal{S}$ that covers $[0,x]$, then it also covers $[0,y] \subseteq [0,x]$. Thus A is an interval.

The set A is bounded above by 1, so it must have a least upper bound; denote it by a. Clearly $0 \leq a \leq 1$. If we can show that

The existence of the least upper bound follows from the fact that we are working in $\mathbb{R}$.

$$a \in A \quad \text{and} \quad a = 1,$$

then the proof will be complete.

Proof that $a \in A$

If $a = 0$, then we know that $a \in A$. So suppose that $a \in (0,1]$. There is a set $U \in \mathcal{S}$ with $a \in U$. Since U is open, we can find r with $0 < r < a$ such that $(a - r, a + r) \subseteq U$. Consider now the point $b = a - \frac{1}{2}r \in (0,1]$ (Figure 4.4). This point certainly lies in A – for otherwise, since A is an interval, each point x between b and a is not in A, contradicting the fact that a is the *least* upper bound of A. Thus the interval $[0,b]$ has a finite subcover $\{U_1, U_2, \ldots, U_n\}$ from $\mathcal{S}$. But then $\{U\} \cup \{U_1, U_2, \ldots, U_n\}$ is a finite subcover from $\mathcal{S}$ of $[0,a]$. It follows that $a \in A$.

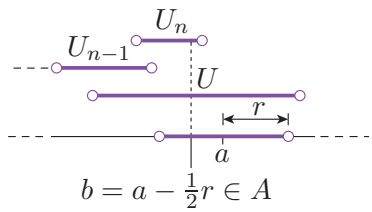

Figure 4.4

Proof that $a = 1$

Suppose, for a contradiction, that $a < 1$. Then, since $a \in A$, the interval $[0,a]$ has a finite subcover $\{U_1, U_2, \ldots, U_n\}$ from $\mathcal{S}$. In particular, $a \in U_i$, for some $i \in \{1, 2, \ldots, n\}$. For this open set U_i, by Book D, Chapter 16, Theorem 2.9 we can find an $r > 0$, with $a + r < 1$ and such that $B_d(a,r) \subseteq U_i$. In particular, $a + \frac{1}{2}r \in U_i$, and so the interval $[0, a + \frac{1}{2}r]$ has a finite subcover from $\mathcal{S}$: namely, $\{U_1, U_2, \ldots, U_n\}$. Thus since $a + \frac{1}{2}r < a + r < 1$, it follows that $a + \frac{1}{2}r \in A$, contradicting the fact that a is an *upper bound* of A. We deduce that $a = 1$.

Note that we use the fact that the U_i are an *open* cover here.

It follows that $\mathcal{S}$ contains a finite subcover of $[0,1]$. ■

Remark

This proof is a classic, created even before formal set theory was invented. An analogous proof can be used to show that any closed bounded interval $[a,b]$ has this property.

Note that it is important that we consider *open* covers of $[0,1]$ – if the cover is not open, there may be no finite subcover.

Exercise 4.4

With respect to the Euclidean metric on $\mathbb{R}$, show that there is an infinite cover of $[0,1]$ by *non-open* sets that has no finite subcover.

We have seen that every open cover of $[0,1]$ has a finite subcover. We now show that the open interval $(0,1)$ does not have this property and ask you to show that the unbounded interval $[0,\infty)$ also does not have this property.

Worked Exercise 4.6

Consider $\mathbb{R}$ with the Euclidean metric and let A be the bounded open set $(0,1)$. Find an open cover of A that has no finite subcover.

Solution

There are many open covers with this property. One such open cover is

$$\mathcal{S} = \{(\tfrac{1}{n}, 1) : n = 2, 3, 4, \ldots\} = \{(\tfrac{1}{2}, 1), (\tfrac{1}{3}, 1), (\tfrac{1}{4}, 1), \ldots\}.$$

For example, if $x = \frac{1}{100}$, then $(\frac{1}{101}, 1)$ contains x.

For each $x \in (0,1)$, the interval $(\frac{1}{n}, 1)$ contains x if $n \in \mathbb{N}$ is chosen so that $\frac{1}{n} < x$. Thus $\mathcal{S}$ is a cover of $(0,1)$. Also, the intervals in $\mathcal{S}$ are open. So $\mathcal{S}$ is an open cover of $(0,1)$. Suppose that $\mathcal{R}$ is a finite subcollection of $\mathcal{S}$. We must show that $\mathcal{R}$ does not cover $(0,1)$. Since $\mathcal{R}$ is finite, we can write

$$\mathcal{R} = \{(\tfrac{1}{n_1}, 1), (\tfrac{1}{n_2}, 1), \ldots, (\tfrac{1}{n_N}, 1)\},$$

where we may assume that $n_1 < n_2 < \cdots < n_N$ (Figure 4.5).

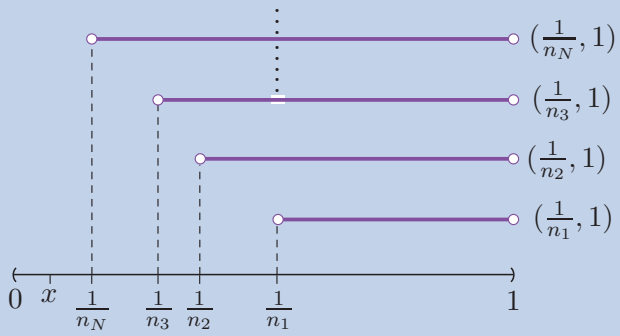

Figure 4.5

Then

$$\bigcup_{U \in \mathcal{R}} U = (\tfrac{1}{n_N}, 1),$$

and so no point $x \in (0, \frac{1}{n_N}]$ is covered. Thus $\mathcal{R}$ is not a cover of $(0, 1)$.

Exercise 4.5

Consider $\mathbb{R}$ with the Euclidean metric and let A be the closed unbounded set $[0, \infty)$. Find an open cover of A that has no finite subcover.

4.1 The Boundedness Theorem revisited

The results of Worked Exercise 4.6 and Exercise 4.5 can be generalised to show that not all $d^{(1)}$-open covers of the open interval (a, b), the half-open intervals $(a, b]$, $[a, b)$, or the unbounded intervals $[a, \infty)$, (a, ∞), $(-\infty, a]$ and $(-\infty, a)$ have finite subcovers. Thus the finite subcover property applies only to intervals that are $d^{(1)}$-closed and bounded.

We now see that the statement 'every open cover of $[0, 1]$ contains a finite subcover' seems to capture some fundamental property of the closed bounded interval $[0, 1]$ that other types of interval fail to possess.

This property enables us to reduce the infinite to the finite. We now give a second proof of the Boundedness Theorem for the interval $[0, 1]$ in which we make use of this ability.

The Boundedness Theorem is Theorem 1.3.

Boundedness Theorem

Let $f : [0, 1] \to \mathbb{R}$ be a continuous function on $[0, 1]$. Then f is bounded on $[0, 1]$. That is, there is a real number M such that $|f(x)| \leq M$ for each $x \in [0, 1]$.

Proof Consider

$$\mathcal{S} = \{U_i : U_i = (i - 1, i + 1),\ i \in \mathbb{Z}\}.$$

Each point $x \in \mathbb{R}$ belongs to an interval U_i. Also, each interval U_i is open. So $\mathcal{S}$ is an open cover of $\mathbb{R}$.

$\mathcal{S}$ is an infinite collection of open sets covering $\mathbb{R}$.

Corresponding to each set U_i is the subset $f^{-1}(U_i)$ of the domain $[0, 1]$. Since f is continuous on $[0, 1]$, the preimage set $f^{-1}(U_i)$ is an open set in $[0, 1]$ (by Book D, Chapter 16, Corollary 7.2). For each $x \in [0, 1]$, $f(x)$ must belong to at least one of the covering sets U_i, so $x \in f^{-1}(U_i)$, and

$$\mathcal{R} = \{f^{-1}(U_i) : i \in \mathbb{Z}\}$$

is an open cover of $[0, 1]$ – that is,

$$[0, 1] \subseteq \bigcup_{i \in \mathbb{Z}} f^{-1}((i - 1, i + 1)).$$

This is where we use compactness to reduce the collection $\mathcal{R} = \{f^{-1}(U_i) : i \in \mathbb{Z}\}$ to the finite collection $\mathcal{R}' = \{f^{-1}(U_i) : -N \leq i \leq N\}$.

Now Theorem 4.5 implies that we can find a finite subcover of $[0, 1]$ from $\mathcal{R}$ – that is, we can find $N \in \mathbb{N}$ such that $[0, 1]$ is covered by the collection $\mathcal{R}' = \{f^{-1}(U_i) : -N \leq i \leq N\}$. (Note that this collection may contain more sets than we need.)

So

$$[0, 1] \subseteq \bigcup_{i=-N}^{N} f^{-1}(U_i).$$

Hence

$$f([0, 1]) \subseteq \bigcup_{i=-N}^{N} U_i = (-N - 1, N + 1).$$

But this means that

$$-N - 1 < f(x) < N + 1 \quad \text{for all } x \in [0, 1].$$

So f is bounded. ■

Thus the finite subcover property of the unit interval is all that is needed to show that a continuous function on that interval is bounded! In particular, note the way in which we used this property to move from an infinite collection of sets, $\bigcup_{i\in\mathbb{Z}} f^{-1}((i - 1, i + 1))$, to a *finite* collection; and that it was this move, from the infinite to the finite, that enabled us to find our bound for the function.

4.2 The definition of a compact metric space

It is also possible to prove the Extreme Value Theorem using the existence of finite subcovers, rather than using the existence of convergent subsequences.

In Subsection 1.1 we saw that every sequence in the closed and bounded interval $[0, 1]$ (with the Euclidean metric) has a convergent subsequence. We used this rather remarkable property in proofs of both the Boundedness Theorem and the Extreme Value Theorem. We also saw, in the discussion after Exercise 1.2, that the unbounded set $\mathbb{R}$ and the open set $(0, 1)$ do not have this property. In Theorem 4.5 we saw that every open cover of the closed and bounded interval $[0, 1]$ (with the Euclidean metric) has a finite subcover. We used this equally remarkable property in a second proof of the Boundedness Theorem (and noted that it could also be used to prove the Extreme Value Theorem). We also saw that unbounded sets, and sets that are not closed, do not possess this property.

In other words, either of these two properties:

1. every sequence has a convergent subsequence
2. every open cover has a finite subcover

can be used to prove results about continuous functions on closed and bounded sets. The first property leads us to a definition of sequential compactness. We now turn our attention to the second, which leads us to the definition of *compactness*.

Definition 4.7

A metric space (X, d) is **compact** if each d-open cover of X contains a finite subcover of X.

A subset $A \subseteq X$ is **compact** if (A, d_A) is compact as a metric space in its own right.

As usual, d_A is the metric on A induced by the metric d on X.

Remark

If we want to emphasise the metric, then we refer to a compact space (X, d) as a d-compact metric space.

We now pause to unpack exactly what this definition means when we are working with subsets, and to check the consistency of our definitions.

Let (X, d) be a metric space with $A \subseteq X$. In Definition 4.4 we defined a d-open cover of A as being a collection of d-open subsets of X which 'cover' A.

There is no problem with applying this definition to the whole set X. We simply take a collection of d-open subsets of X which cover all of X. We do need to be a little careful when we apply this definition to a compact *subset* of X.

Definition 4.7 states that A is a compact subset of X if (A, d_A) is compact as a metric space in its own right. That is, A is a compact subset if every d_A-open cover of A has a finite subcover. We know, from Theorem 2.1 of Chapter 21, that a set $U \subseteq A$ is d_A-open in A if, and only if, there is a d-open set $V \subseteq X$ with $U = A \cap V$. Therefore, for every d-open cover of A, by d-open sets, $V_i \subseteq X$, we have a corresponding d_A-open cover of A, given by the collection of all sets of the form $V_i \cap A$. Hence, if an open cover of A by d-open sets in X has a finite subcover, then the corresponding open cover by d_A-open sets in A also has a finite subcover, and vice versa; see Figure 4.6. Therefore our definitions are consistent.

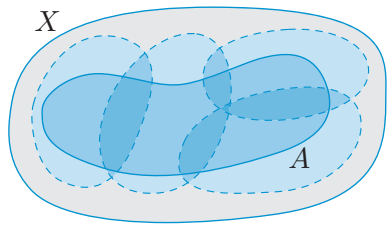

V_i covering $A \subseteq X$

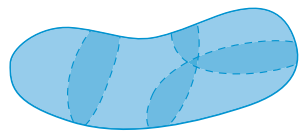

$V_i \cap A$ covering A

Figure 4.6

Corollary 4.8

The interval $[0, 1]$, with the metric induced by the usual Euclidean metric on $\mathbb{R}$, is compact.

Proof This follows from Theorem 4.5. ■

A similar argument can be used to show that each $d^{(1)}$-closed and bounded interval $[a, b]$ is compact.

$(0, 1)$ and $[0, \infty)$ are not compact, as is the case for any interval that is not both closed and bounded.

We state this formally in Theorem 5.14

Remark

Warning: when trying to prove that a given space X is compact, it is tempting, but wrong, to choose some *particular* open cover and demonstrate that it has a finite subcover. This does not prove that X is

compact. We must show that *every* open cover has a finite subcover. Thus a proof of compactness typically begins by considering an *arbitrary* open cover $\mathcal{S}$ of X. To show that a space is not compact it is, of course, sufficient to find a single open cover which does not have a finite subcover.

4.3 Examples of compact spaces

We now ask you to examine some metric spaces to see which are compact.

Exercise 4.6

Let (X, d) be a metric space. Show that $\emptyset$ is a compact subset of X.

The most important examples for us to examine for compactness are subsets of $\mathbb{R}^n$. Ideally we would like to be able to characterise these completely – that is we would like necessary and sufficient conditions for a subset of $\mathbb{R}^n$ to be compact. We already know, from Proposition 2.11, that any $d^{(n)}$-closed and bounded subset of $\mathbb{R}^n$ (with the usual Euclidean metric) is sequentially compact. We would like to use this fact to help us to characterise compact subsets of $\mathbb{R}^n$, so we defer discussion of these subsets until Section 5, where we show that any sequentially compact metric space is compact.

For the moment, we first consider finite sets. It is no surprise to discover that these are compact.

Exercise 4.7

Let (X, d) be a metric space. Show that each finite subset of X is compact.

We now turn our attention to the discrete metric.

Example 4.9

Here every subset of X is open.

Consider (X, d_0), where d_0 is the discrete metric. We know from Exercise 4.7 that any finite set is compact, so suppose that X is an infinite set and let A be an infinite subset of X. In this case, A is not compact. To show this, it is enough to exhibit one open cover of A that has no finite subcover: the collection $\{\{a\} : a \in A\}$ is such a cover.

Exercise 4.8

Show that $(\mathbb{R}, d^{(1)})$ is not compact.

5 Sequentially compact and compact are equivalent for metric spaces

We divide the proof that sequential compactness and compactness are equivalent notions for metric spaces into two parts. First we show in Theorem 5.10 that any sequentially compact metric space is compact, and then in Theorem 5.11 we show that compact metric spaces are sequentially compact.

The main idea of the proof of Theorem 5.10 is to show that all the points in a sequentially compact space are 'near' one of the points in a finite set of points. This allows us to 'replace' the space by a finite collection of open sets. As with many proofs preceded by a discussion where some of the words are held within quotation marks, the proof that a sequentially compact metric space is compact is rather long.

We want to show that all the points in a sequentially compact metric space, (X, d), can be approximated by a finite set of points. By this we mean that for each $\varepsilon > 0$, it is possible to find a finite set $N(\varepsilon)$ of points, such that for each point $x \in X$, we can find a $p_x \in N(\varepsilon)$, with $d(x, p_x) < \varepsilon$. This allows us to 'replace' the whole space X by the finite set of points, $N(\varepsilon)$, with an error of at most ε.

The set $N(\varepsilon)$ will depend on ε and the point p_x will also depend on x.

Definition 5.1 *ε-net*

Let (X, d) be a metric space. For $\varepsilon > 0$, an **ε-net** for X is a set $N(\varepsilon) \subseteq X$ such that

$$X \subseteq \bigcup_{p \in N(\varepsilon)} B_d(p, \varepsilon).$$

Remark

When we want to emphasise the metric, we use the notation $N_d(\varepsilon)$.

An ε-net for a metric space (X, d) is a set of points, $\{p_i : i \in I\}$, such that taking a d-open ball, of radius ε, around each of these points, gives us a collection of balls that covers the space X (see Figure 5.1 overleaf).

Remarks

1. Note that an ε-net is a set of *points*. It does not look like a fishing net (or for that matter a net of any other description). It is a (possibly uncountable) set of points such that any other point in the space is 'close' to one of the points in the net.
2. An ε-net provides an approximation to within a distance ε of any point in the space in the following sense. If $N(\varepsilon) = \{p_i : i \in I\}$ is an ε-net for the metric space (X, d), then for each $x \in X$, we can find a $p_i \in N(\varepsilon)$ such that $d(x, p_i) < \varepsilon$. For example, $\mathbb{Z}$ is an ε-net for $\mathbb{R}$ with the Euclidean metric, provided that $\varepsilon > \frac{1}{2}$.

3. For each $\varepsilon > 0$, X is an ε-net for itself.
4. If S is an ε-net for X, then it is also a δ-net for X for each $\delta > \varepsilon$.

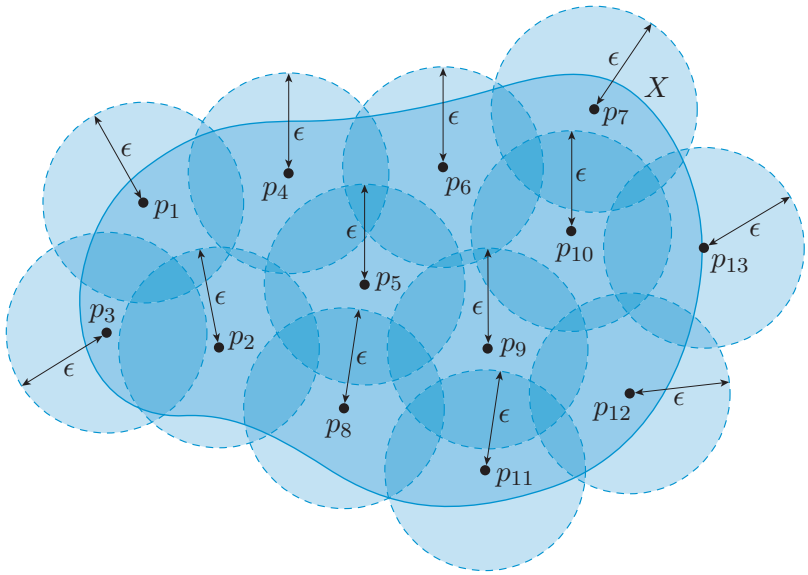

Figure 5.1

Every metric space, (X, d), has at least one ε-net for any ε, namely X itself. (Of course, if X is infinite, then this will be an infinite ε-net.) As the first step of our journey to showing that every sequentially compact space is compact, we show that a sequentially compact space has a *finite* ε-net for each $\varepsilon > 0$. This will not, of course, show that every open cover has a finite subcover; but it will give us a finite collection of sets of radius ε with which to work.

Lemma 5.2

Let (X, d) be a sequentially compact metric space. Then X has a *finite* ε-net for each $\varepsilon > 0$.

Proof Clearly the empty set has a finite ε-net for any ε, so we may assume that $X \neq \emptyset$. Suppose that X is sequentially compact, but that there is an $\varepsilon > 0$ for which there is no finite ε-net for X. We obtain a contradiction by constructing a sequence in X that has no convergent subsequence. We will do this by taking an arbitrary point, $a_1 \in X$, and then taking a second point which is 'more than ε distant' from a_1, then finding a third which is 'more than ε distant' from each of these, and so on, which gives us a sequence where all of the points are spaced at least ε apart (and so it cannot have a convergent subsequence).

Let $a_1 \in X$. Since there is no finite ε-net, $\{a_1\}$ cannot be an ε-net for X. Hence there is an $a_2 \in X$ such that $d(a_1, a_2) \geq \varepsilon$.

Similarly, $\{a_1, a_2\}$ cannot be an ε-net for X, and so there is an $a_3 \in X$ whose distance is at least ε from both a_1 and a_2 – that is,

$$\min\{d(a_1, a_3), d(a_2, a_3)\} \geq \varepsilon.$$

Thus for $\{a_1, a_2, a_3\}$, we have $d(a_i, a_j) \geq \varepsilon$ for each $i \neq j$.

We continue in this way, constructing subsets $\{a_1, a_2, \ldots, a_k\}$ of X that are not ε-nets for X and for which $d(a_i, a_j) \geq \varepsilon$ for $i \neq j$.

Thus we construct a sequence (a_n) such that $d(a_i, a_j) \geq \varepsilon$ for any two distinct terms a_i and a_j in the sequence. In particular, the terms of this sequence do not get arbitrarily close (they are always at least ε apart), and so this sequence has no convergent subsequences, contradicting our assumption that X is sequentially compact. ■

Exercise 5.1

Use Lemma 5.2 to show that a sequentially compact metric space is bounded.

We have proved the following.

Corollary 5.3

A sequentially compact metric space is bounded.

Proof This is Exercise 5.1. ■

We have now seen that a sequentially compact metric space has a finite ε-net, and that this fact can be used to show that any sequentially compact metric space is bounded. However, notice that we can prove that a sequentially compact space is bounded by assuming that a finite ε-net exists *for just one value of* ε. Lemma 5.2 shows something far stronger: that, in a sequentially compact metric space, we can find a finite ε-net for *any* $\varepsilon > 0$.

5.1 Totally bounded sets

We want to show that sequential compactness is equivalent to compactness for metric spaces. We have just shown, in Corollary 5.3, that all sequentially compact metric spaces are bounded. This may lead you to wonder if sequential compactness is equivalent to boundedness. This turns out to be far from the case. In fact, the next few exercises show that, when investigated using only convergent sequences, all metric spaces 'look bounded'! To do this we introduce a new metric, called the **min metric**: $d_{\min}$.

Definition 5.4 *The min metric*

Let (X, d) be a metric space and define a distance function $d_{\min}\colon X \times X \to \mathbb{R}$ as follows:

$$d_{\min}(a,b) = \begin{cases} d(a,b), & \text{if } d(a,b) \leq 1, \\ 1, & \text{if } d(a,b) > 1. \end{cases}$$

First we ask you to check that $d_{\min}$ really is a metric (assuming that d is a metric to start with, of course).

Exercise 5.2

Show that $d_{\min}$ is a metric.

The following may help you to develop a feel for a space with the min metric. Consider the real line, with the metric $d^{(1)}_{\min}$. Imagine that we are standing in this space, at a fixed point, and looking around. Intuitively, when we 'look at' our immediate surroundings, we can see some detail (as points that are 'within 1 of us' retain their original spacings), but once we look further into the distance all of the points appear equally far away. We can easily bound our space by enclosing it within a ball of any radius greater than 1.

We now show that the convergent sequences in any metric space (X, d) are the same as the convergent sequences in the space with the corresponding min metric, $d_{\min}$.

Exercise 5.3

Let (X, d) be a metric space and let (a_n) be a sequence in X. Show that (a_n) is d-convergent if, and only if, (a_n) is $d_{\min}$-convergent.

Exercise 5.3 shows that the convergent sequences for any metric space (X, d) are the same as those for the metric space $(X, d_{\min})$. We now show that any subset of a space with the min metric is $d_{\min}$-bounded.

Exercise 5.4

Let (X, d) be any metric space and let $A \subseteq X$. Show that A is $d_{\min}$-bounded, where $d_{\min}$ is the metric derived from d as described in Definition 5.4.

In particular, the whole space $(X, d_{\min})$ is bounded. Therefore the convergent sequences in any metric space (X, d) are the same as the convergence sequences in the metric space $(X, d_{\min})$, and in this space every set is bounded. So, if we only use convergent sequences to examine a metric space, then any space (X, d) will 'look like' the bounded space

$(X, d_{\min})$. In other words, if we only use sequences to 'look with', then we cannot 'see' unboundedness. Since the definition of sequential compactness only relies on the convergence of certain (sub)sequences, sequential compactness is a property that we need to use sequences to 'see'. Therefore boundedness is not going to be a helpful quality when trying to prove that sequentially compact spaces are compact and we need a stronger concept than boundedness. This stronger concept is that of *total boundedness*, the definition of which is motivated by Lemma 5.2.

Definition 5.5 *Totally bounded*

A metric space (X, d) is **totally bounded** if there is a finite ε-net for each $\varepsilon > 0$.

In a bounded space, we can always find an $\varepsilon > 0$ for which there is a finite ε-net. In a totally bounded space, rather than being satisfied by finding a single ε, and corresponding finite ε-net, we require that we can start with an arbitrary $\varepsilon > 0$ and, no matter how small this ε is, guarantee that we can always find a finite ε-net for this ε. In general, as ε gets smaller, the number of points in the ε-net will get larger, but for a totally bounded space this number must remain finite, however small ε becomes. In the next exercise we ask you to show that the closed and bounded interval $[a, b]$ is, in fact, closed and totally bounded.

Exercise 5.5

Show that $[a, b]$ (with the Euclidean metric) is totally bounded.

We now return very briefly to the min metric. Consider $(\mathbb{R}, d_{\min}^{(1)})$, which we know is bounded. For this space to be totally bounded, it is not sufficient to fix a point and simply enclose our space within a finite-radius ball, where the radius can be as big as we like; we need to be able to find finitely many balls for any given radius ε. Taking $\varepsilon < 1$ means that we are unable to take advantage of the 'contracting effect' of the min metric and so would need infinitely many balls to cover our space. So we see that total boundedness is a stronger concept than boundedness.

In the next exercise, we ask you to show that every compact metric space is totally bounded.

Exercise 5.6

Show directly that each compact metric space (X, d) is totally bounded.

5.2 Sequentially compact metric spaces are compact

In our quest to prove that sequential compactness and compactness are equivalent for metric spaces, we now have the following two results.

Lemma 5.6

Each compact metric space (X, d) is totally bounded.

Proof This is Exercise 5.6. ■

Corollary 5.7

Each sequentially compact metric space is totally bounded.

Proof This follows directly from Lemma 5.2 and Definition 5.5. ■

The situation is looking hopeful. We now show that, for any open cover $\mathcal{S}$ of a sequentially compact space (X, d), we can find a number ε such that the d-open ball of radius ε around each point lies entirely within one of the sets in $\mathcal{S}$. This may seem a rather dull and technical result, but it turns out to be the key to proving that sequentially compact metric spaces are compact.

Lemma 5.8

Let (X, d) be a sequentially compact metric space and let $\mathcal{S}$ be an open cover of X. Then there is an $\varepsilon > 0$ such that, for each $x \in X$, there is a $U \in \mathcal{S}$ with $B_d(x, \varepsilon) \subseteq U$.

Proof Let (X, d) be a sequentially compact metric space, and suppose that the result is false.

In this case, there is an open cover $\mathcal{S}$ of X such that we cannot find an $\varepsilon > 0$ such that, for each $x \in X$, there is a $U \in \mathcal{S}$ with $B_d(x, \varepsilon) \subseteq U$. Therefore for each $\varepsilon > 0$, there must be an $x \in X$ with

This says that the ball $B_d(x, \varepsilon)$ is never entirely contained within an open set of the cover.

$$B_d(x, \varepsilon) - U \neq \emptyset \quad \text{for each } U \in \mathcal{S}.$$

We now use this fact to construct a sequence.

We are taking $\varepsilon = \frac{1}{n}$ here.

For each $n \in \mathbb{N}$, there is an $a_n \in X$ such that

$$B_d(a_n, \tfrac{1}{n}) - U \neq \emptyset \quad \text{for each } U \in \mathcal{S}.$$

Since X is sequentially compact, the sequence $(a_n)_{n=1}^{\infty}$ has a convergent subsequence $(a_{n_k})_{k=1}^{\infty}$ with limit $a \in X$, say.

Since $\mathcal{S}$ is a cover of X, there is a set $U \in \mathcal{S}$ with $a \in U$. Since U is open, there is an $M \in \mathbb{N}$ such that

$$B_d\left(a, \tfrac{2}{M}\right) \subseteq U.$$

Now, $a_{n_k} \to a$ as $k \to \infty$, and so there is a $K \in \mathbb{N}$ such that, for all $k > K$,

$$d(a_{n_k}, a) < \tfrac{1}{M}, \quad \text{and so } a_{n_k} \in B_d\left(a, \tfrac{1}{M}\right).$$

We now recall that, by the definition of the sequence (a_n),

$$B_d\left(a_{n_k}, \tfrac{1}{n_k}\right) - U \neq \emptyset \quad \text{for each } k \in \mathbb{N}.$$

However, if k is chosen so that both $k > K$ and $n_k > M$, then

$$B_d\left(a_{n_k}, \tfrac{1}{n_k}\right) \subseteq B_d\left(a_{n_k}, \tfrac{1}{M}\right) \subseteq B_d\left(a, \tfrac{2}{M}\right) \subseteq U.$$

See Figure 4.1.

This implies that $B_d\left(a_{n_k}, \tfrac{1}{n_k}\right) - U = \emptyset$: a contradiction.

Hence there is an $\varepsilon > 0$ such that, for each $x \in X$, there is a $U \in \mathcal{S}$ with $B_d(x, \varepsilon) \subseteq U$. ■

Definition 5.9 *Lebesgue number*

The number ε given in Lemma 5.8 is called a **Lebesgue number** of the open cover $\mathcal{S}$.

Henri Lebesgue (1875–1941) is primarily remembered by mathematicians for his fundamental contributions to the theory of integration.

We can now prove that sequentially compact metric spaces are compact.

Theorem 5.10

Every sequentially compact metric space is compact.

Proof Let (X, d) be a sequentially compact metric space. We must show that each open cover of X has a finite subcover.

So let $\mathcal{S}$ be an open cover of X. Lemma 5.8 implies that there is an $\varepsilon > 0$ such that, for each $x \in X$, there is a $U \in \mathcal{S}$ with $B_d(x, \varepsilon) \subseteq U$.

Since (X, d) is sequentially compact, Corollary 5.7 implies that (X, d) is totally bounded. Thus there are a finite number of points $x_1, x_2, \ldots, x_N \in X$ such that, for this choice of ε,

$$X \subseteq \bigcup_{i=1}^{N} B_d(x_i, \varepsilon).$$

But, for each x_i, there is a $U_i \in \mathcal{S}$ for which $B_d(x_i, \varepsilon) \subseteq U_i$. Therefore

$$X \subseteq \bigcup_{i=1}^{N} B_d(x_i, \varepsilon) \subseteq \bigcup_{i=1}^{N} U_i.$$

Thus the collection $\{U_1, U_2, \ldots, U_N\}$ is a finite subcover of X from $\mathcal{S}$.

Since $\mathcal{S}$ is an arbitrary open cover of X, the result follows. ■

5.3 Compact metric spaces are sequentially compact

In Theorem 5.10 we proved that all sequentially compact metric spaces are compact. We now show that all compact metric spaces are sequentially compact, thus proving that, for metric spaces, the definitions are equivalent.

Theorem 5.11

Let (X, d) be a compact metric space. Then each sequence in X has a convergent subsequence and so (X, d) is sequentially compact.

Proof The proof is by contradiction, twice.

Suppose that (X, d) is a compact metric space for which there is a sequence (a_n) in X with no convergent subsequences.

We claim that, under this supposition, for each $x \in X$, there are an $\varepsilon_x > 0$ and an $N_x \in \mathbb{N}$ such that

$$a_n \notin B_d(x, \varepsilon_x) \quad \text{for all } n > N_x.$$

We prove this claim by contradiction. Suppose that the claim is not true. Then there is an $x \in X$ such that, for each $\varepsilon > 0$ and each $N \in \mathbb{N}$, there is some $n > N$ for which

$$a_n \in B_d(x, \varepsilon).$$

In particular, for $\varepsilon_1 = \frac{1}{2}$, there is an integer n_1 such that $n_1 > 1$ and

$$a_{n_1} \in B_d(x, \varepsilon_1) = B_d(x, \tfrac{1}{2}).$$

Similarly, for $\varepsilon_2 = \frac{1}{4}$, there is an integer $n_2 > n_1$ such that

$$a_{n_2} \in B_d(x, \varepsilon_2) = B_d(x, \tfrac{1}{4}).$$

Proceeding in this way, we find that for $\varepsilon_k = 2^{-k}$ there is an integer n_k such that $n_k > n_{k-1} > \cdots > n_1$ and

$$a_{n_k} \in B_d(x, \varepsilon_k) = B_d(x, 2^{-k}).$$

This follows from the Squeeze Rule as (2^{-k}) is a null sequence.

Therefore $d(a_{n_k}, x) < 2^{-k}$ and so $(d(a_{n_k}, x))$ is a null sequence and we have a subsequence (a_{n_k}) that converges to x. But we are assuming that no such subsequence exists – a contradiction!

Thus, under the supposition that the sequence (a_n) has no convergent subsequences, we have shown that, for each $x \in X$, there are an $\varepsilon_x > 0$ and an $N_x \in \mathbb{N}$ with

$$B_d(x, \varepsilon_x) \cap \{a_n : n > N_x\} = \emptyset.$$

Clearly the collection

$$\{B_d(x, \varepsilon_x) : x \in X\}$$

This is where we use compactness.

is an open cover of X. Since X is compact, the definition of compactness tells us that there exist $x_1, x_2, \ldots, x_k \in X$ such that

$$X \subseteq \bigcup_{i=1}^{k} B_d(x_i, \varepsilon_{x_i}).$$

In particular, the set $\{a_n : n \in \mathbb{N}\}$ of the sequence is contained in this finite union of balls:

$$\{a_n : n \in \mathbb{N}\} \subseteq \bigcup_{i=1}^{k} B_d(x_i, \varepsilon_{x_i}).$$

By the definition of N_{x_i},

$$B_d(x_i, \varepsilon_{x_i}) \cap \{a_n : n > N_{x_i}\} = \emptyset.$$

Now let $N = \max\{N_{x_1}, N_{x_2}, \ldots, N_{x_k}\}$. It follows that

$$a_n \notin \bigcup_{i=1}^{k} B_d(x_i, \varepsilon_{x_i}) \quad \text{for } n > N$$

– a contradiction, since the whole space is contained in this union!

We conclude that there is no sequence in X without a convergent subsequence. ■

Thus we have:

Theorem 5.12

A metric space is compact if, and only if, it is sequentially compact.

Proof This follows from Theorem 5.11 and Theorem 5.10. ■

Remark

Note that if (X, d) is a metric space and (a_n) lies in a compact subset A of X, then (a_n) has a convergent subsequence with limit in A.

We now know that, for metric spaces, the notions of compact and sequentially compact are equivalent. Therefore our results for sequentially compact spaces hold for compact spaces and we have the following results.

Corollary 5.13

1. Let $A \subseteq X$ where (X, d) is a d-compact metric space. Then A is d-closed if, and only if, A is d_A-compact.
2. Let (X, d) and (Y, d') be metric spaces with X being d-compact. Let $f\colon X \to Y$ be a (d, d')-continuous function. Then the image, $f(X)$, is d'-compact.
3. Let (X, d) and (Y, d') be two compact metric spaces. Then $(X \times Y, e)$ is compact (where e is a product metric).
4. Compactness is a topological invariant.

Proof

1. This follows from Theorem 2.5 and Theorem 5.12.
2. This follows from Lemma 2.6 and Theorem 5.12.
3. This follows from Theorem 2.8 and Theorem 5.12.
4. This follows from Corollary 2.7 and Theorem 5.12.

■

In particular, we have the following version of the Heine–Borel Theorem.

Theorem 5.14 *The Heine–Borel Theorem*

A subset of $\mathbb{R}^n$ is $d^{(n)}$-compact if, and only if, it is $d^{(n)}$-closed and bounded.

Proof This follows from Theorem 2.12 and Theorem 5.12. ■

6 Unions and intersections of compact spaces

Since compactness is expressed in terms of open sets, it can be easier to use compactness (as opposed to sequential compactness) to prove results concerning set operations.

Theorem 6.1

Let (X, d) be a metric space. Then the intersection of any collection of compact subsets of X is compact.

Proof Let $\mathcal{F}$ be a collection of d-compact subsets of X. We must show that $\bigcap_{K\in\mathcal{F}} K$ is d-compact. Theorem 5.12 implies that K is d-sequentially compact and so, by the result of Exercise 2.1, K is d-closed. It follows that $\bigcap_{K\in\mathcal{F}} K$ is a d_K-closed subset (for the induced metric) of any one of the compact sets K in the collection $\mathcal{F}$, and so, by Corollary 5.13, $\bigcap_{K\in\mathcal{F}} K$ is compact. ■

Corollary 6.2

Let (X, d) be a metric space. Then the intersection of any collection of sequentially compact subsets of X is sequentially compact.

Proof This follows from Theorem 6.1 and Theorem 5.12. ■

In Chapter 24, we will investigate some basic properties of fractal sets. You will see that many such sets can be defined as intersections of compact subsets of metric spaces. Theorem 6.1 guarantees that the resulting intersection is compact. However, the intersection can be empty, particularly if a collection $\mathcal{F}$ of compact sets is large. For example, if $\mathcal{F} = \{\{x\} : x \in [0,1]\}$, then $\bigcap_{K\in\mathcal{F}} K = \emptyset$. Fortunately, there is one situation, useful in the study of fractals, where the intersection of a collection of compact sets must be non-empty in a metric space.

Theorem 6.3

Let (X, d) be a metric space, and let K_n $(n = 1, 2, 3, \ldots)$ be non-empty compact sets in X for which

$$K_1 \supseteq K_2 \supseteq K_3 \supseteq \cdots \supseteq K_n \supseteq K_{n+1} \supseteq \cdots .$$

Then $\bigcap_{n=1}^{\infty} K_n$ is a non-empty compact set.

Proof Theorem 6.1 implies that $\bigcap_{n=1}^{\infty} K_n$ is compact. It remains to show that the intersection is non-empty.

Since K_i is compact, for each i, it follows from Corollary 5.13 that each K_n is closed. Thus K_n^c is open, for each n. Moreover, by De Morgan's Second Law,

Chapter 16, Theorem 2.12.

$$\bigcup_{n=1}^{\infty} K_n^c = \left(\bigcap_{n=1}^{\infty} K_n\right)^c .$$

We assume that $\bigcap_{n=1}^{\infty} K_n = \emptyset$ and use proof by contradiction. Then $\{K_n^c : n \in \mathbb{N}\}$ is an open cover of X, and so (in particular) is an open cover of K_1. Since K_1 is compact, there is a number $N \in \mathbb{N}$ for which

$$K_1 \subseteq \bigcup_{n=1}^{N} K_n^c \quad \text{(Figure 6.1)}.$$

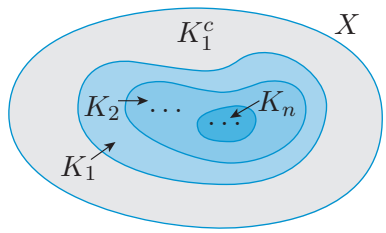

Figure 6.1

Using De Morgan's Second Law again, we see that

$$\bigcup_{n=1}^{N} K_n^c = \left(\bigcap_{n=1}^{N} K_n\right)^c = K_N^c.$$

Therefore $K_1 \subseteq K_N^c$. Hence

$$K_N \subseteq K_1 \subseteq K_N^c.$$

We deduce that K_N is empty – a contradiction. Hence $\bigcap_{n=1}^{\infty} K_n \neq \emptyset$, as required. ■

Proposition 6.4

The finite union of a collection of compact subsets is compact.

Proof We ask you to prove this in Exercise 6.1. ■

Corollary 6.5

The finite union of a collection of sequentially compact subsets is sequentially compact.

Proof This follows from Proposition 6.4 and Theorem 5.12. ■

Exercise 6.1

Let (X, d) be a metric space, and let K_m $(m = 1, \ldots, n)$ be non-empty compact sets in X. Show that the finite union, $\bigcup_{m=1}^{n} K_m$ is a non-empty compact set.

6.1 Epilogue

This section is not assessed.

You may be wondering why we have made a point of adding 'for a metric space' to our results and comments about the equivalence of the definitions of sequential compactness and compactness. This is because the two definitions can be made to apply to all topological spaces, not just to topological spaces which are also metric spaces (metrisable topological spaces) and, for a topological space, the two definitions are not equivalent. A metric space, (X, d), can be viewed as a topological space by taking the d-open sets as its topology, but it is also possible to specify topologies which do not arise from any metric. In a metrisable topological space (one where the open sets arise as the d-open sets for metric d), the notions of sequential compactness and compactness coincide; but in a general topological space this is not true. For example, it turns out that there is a topological space whose underlying set consists of uncountably many copies of the two-point space $\{0, 1\}$ that is compact but not sequentially compact; and there are topological spaces which are sequentially compact but not compact. However, a discussion of these spaces, and the way in which we can define topologies on them, is beyond the scope of this module.

A topological space was defined in Book D, Chapter 16 and that the structure on a topological space is given by specifying a collection of sets, which we call 'open'. This collection is called a *topology*.

Learning outcomes

We list those topics on which we may set assessment questions to test your understanding of the chapter.

After working through this chapter you should be able to:

(a) Explain what it means for a metric space to be sequentially compact.

(b) Identify sequentially compact subsets of many of our standard examples of metric spaces.

(c) Use the Heine–Borel Theorem to identify sequentially compact subsets of Euclidean spaces.

(d) Use the General Extreme Value Theorem.

(e) Know what it means for a collection of functions to be pointwise bounded or equicontinuous.

(f) Use the Arzelà–Ascoli Theorem to identify sequentially compact subsets of $C[0, 1]$.

(g) Explain what it means for a metric space to be compact.

(h) Identify compact subsets of many metric spaces.

(i) Know that compactness and sequential compactness are equivalent in metric spaces.

(j) Know how compact and sequentially compact sets behave under unions and intersections.

Solutions to exercises

Solution to Exercise 1.1

In an arbitrary metric space we have a notion of distance but we do not necessarily have a notion of *order*. In other words, $a > b$ does not automatically make sense. For example, the complex numbers $\mathbb{C}$ with the metric given by $d(w, z) = |w - z|$ where $|a + bi| = \sqrt{a^2 + b^2}$ form a metric space, but it is not obvious which of, for example, $2 + i$ and $1 + 2i$ is, or should be, 'bigger'.

Solution to Exercise 1.2

Points in $C[0, 1]$ are functions, so a sequence in $C[0, 1]$ will be an ordered list of functions, $(f_1, f_2, f_3, \ldots)$, where, for each $n \in \mathbb{N}$, f_n is a continuous function on $[0, 1]$. Therefore the following is an example of a sequence in $C[0, 1]$. Note that this is just one example. There are (infinitely) many more.

For each n, define $f_n \colon [0, 1] \to \mathbb{R}$ by $f_n(x) = nx$. This gives the sequence
$f_1 \colon x \to x$
$f_2 \colon x \to 2x$
$f_3 \colon x \to 3x$
$\vdots$

Solution to Exercise 1.3

A point in Cantor space $\mathbf{C}$, is an infinite sequence of zeros and ones. So a *sequence* in Cantor space is an ordered list of sequences (in other words, a 'sequence of sequences'). Therefore the following is an example of a sequence in Cantor space. Note that this is just an example. There are (infinitely) many more:
$(0, 0, 1, 1, 0, 0, 1, \ldots)$
$(0, 1, 0, 1, 0, 1, 0, \ldots)$
$(0, 0, 0, 1, 1, 1, 0, \ldots)$
$\vdots$

Solution to Exercise 1.4

An example is the sequence $\left(\frac{1}{3^n}\right)$ (since $\frac{1}{3^n} \in K_C$ for all n and this sequence is $d^{(1)}_{K_C}$-convergent to 0).

Solution to Exercise 1.5

Assume that we have chosen the first two terms a_{N_1} and a_{N_2} as described in the proof of Lemma 1.1.

Let $P(m)$ be the proposition that we can find m terms $a_{N_1}, \ldots, a_{N_m}$ in (a_n) with $a_{N_j} \geq a_{N_i}$ for all $j \geq i$. Then $P(2)$ is true, so we have a basis for the induction. Assume that $P(k)$ is true. We know that a_{N_k} is not a peak term of the sequence (a_n) and so there must be a term a_l with $l > k$ and $a_l \geq a_k$. Set $a_{N_{k+1}} = a_l$. Then, since $a_l \geq a_k$, by the induction hypothesis we have that $a_{N_1}, \ldots, a_{N_{k+1}}$ are $k+1$ terms with $a_{N_j} \geq a_{N_i}$ for all $j \geq i$ which is $P(k+1)$. Hence, by mathematical induction, we have the result.

Solution to Exercise 2.1

Let $A \subseteq X$ be a sequentially compact subset of a metric space (X, d). To show that A is d-closed, we need to show that any d-convergent sequence with terms in A has a limit in A. Let (a_n) be a d-convergent sequence in A, with limit a. Since A is sequentially compact, (a_n) has a d_A-convergent subsequence, (a_{n_i}), with a limit, a', in A. By Book D, Chapter 15, Theorem 1.5, we have that (a_{n_i}) is d-convergent to a', and so by Book D, Chapter 16, Lemma 1.13 we have $a' = a$ and $a \in A$. Therefore A is d-closed.

Solution to Exercise 2.2

Since A is not bounded, we cannot find a real number M such that $|x| \leq M$ for each $x \in A$. Therefore for each $n \in \mathbb{N}$ we can find an element $a_n \in A$ such that $|a_n| \geq n$. This gives a sequence (a_n). Assume (for a contradiction) that we can find a convergent subsequence (a_{n_i}) with (a_{n_i}) converging to a limit L. In other words, for each $\varepsilon > 0$, there is an $N \in \mathbb{N}$ for which $d(a_{n_i}, L) < \varepsilon$ for all $i > N$. Set $\varepsilon = 1$.

See the second remark after Book D, Chapter 14, Definition 3.2.

There is nothing special about the choice of $\varepsilon = 1$ here. We just need to fix the size of a 'gap' between the possible limit and the 'tail' of the unbounded sequence.

Now $n_i \geq i$ (since (a_{n_i}) is a subsequence of (a_n)), and so for each i we have $|a_{n_i}| \geq n_i \geq i$. Pick $M > |L| + 1$. Then $|a_{n_i}| > |L| + 1$ for all $i \geq M$, and so $||a_{n_i}| - |L|| > 1$ for all $i \geq M$. Therefore, for all $i \geq M$ we have $d^{(1)}(a_{n_i}, L) = |a_{n_i} - L| \geq ||a_{n_i}| - |L|| > 1$ (by the Reverse Triangle Inequality). Therefore it is not possible to find any $N \in \mathbb{N}$ for which $d(a_{n_i}, L) < 1$ for all $i > N$, which contradicts that fact that (a_{n_i}) converges to the limit L.

Solution to Exercise 2.3

The middle-third Cantor set $K_C = \bigcap_{n=1}^{\infty} K_n$ where:

$$K_0 = [0,1]$$
$$K_1 = \left[0, \tfrac{1}{3}\right] \cup \left[\tfrac{2}{3}, 1\right]$$
$$K_2 = \left[0, \tfrac{1}{9}\right] \cup \left[\tfrac{2}{9}, \tfrac{1}{3}\right] \cup \left[\tfrac{2}{3}, \tfrac{7}{9}\right] \cup \left[\tfrac{8}{9}, 1\right]$$
$$\vdots$$

with K_n formed by removing the open middle third of each interval in K_{n-1}. Consider K_i for any i. Then K_i is a finite union of closed intervals and so, by the paragraph below Book D, Chapter 16, Proposition 1.14, K_i is closed for any i. Now $K_C = \bigcap_{n=1}^{\infty} K_n$ and so is the infinite intersection of a collection of closed sets, and so is closed by Proposition 1.17 of Chapter 16. Also $K_C \subseteq [0,1]$, and so $|c| \leq 1$ for all $c \in K_C$, and K_C is bounded. Therefore by Theorem 1.2 K_C is sequentially compact.

Solution to Exercise 2.4

Let (a_n) be a sequence in X. As (a_n) contains infinitely many terms, then at least one element of X, say x, must appear infinitely many times in the sequence. Therefore we must have $a_{n_i} = x$ for infinitely many n_i, giving a constant subsequence which is d-convergent to x (by Book D, Chapter 14, Worked Exercise 3.5).

Solution to Exercise 2.5

Let (f_n) be the sequence of functions in A given by

$$f_n(x) = n \quad \text{for each } x \in [0,1].$$

For each $x \in [0,1]$, the sequence $(f_n(x))$ diverges, and so has no pointwise limit; moreover, any subsequence of $(f_n(x))$ also diverges. Lemma 4.5 of Chapter 15 implies that there is no uniformly convergent subsequence. But, by Theorem 5.8 of Chapter 15, uniform convergence is the same as $d_{\max}$-convergence in $C[0,1]$. It follows that there are no $d_{\max}$-convergent subsequences of (f_n). Hence A is not sequentially compact.

Solution to Exercise 2.6

Let (X,d) and (Y,e) be metric spaces with X being sequentially compact. Let $f\colon X \to Y$ be a (d,e)-continuous function and let (a_n) be a sequence in the image $f(X)$. We need to find an $e_{f(X)}$-convergent subsequence. Since (a_n) is in $f(X)$, we can find a sequence (x_n) such that $f(x_n) = a_n$ for each n. Then (x_n), being a sequence in the sequentially compact space X, has a convergent subsequence (x_{n_i}), converging to a limit, say x, in X. Then, since f is continuous, we must have that $\big(f(x_{n_i})\big)$ converges to $f(x) \in f(X)$, and so $\big(f(x_{n_i})\big) = (a_{n_i})$ is a convergent subsequence of $(f(x_n)) = (a_n)$.

Solution to Exercise 2.7

Let M be such that $d^{(n)}(\mathbf{x}, \mathbf{y}) \leq M$ for all $\mathbf{x}, \mathbf{y} \in A$. Let $\mathbf{a} = (a_1, \ldots, a_n)$, $\mathbf{b} = (b_1, \ldots, b_n)$ be in A. Since $d^{(n)}(\mathbf{a}, \mathbf{b}) = \sqrt{(b_1 - a_1)^2 + (b_2 - a_2)^2 + \cdots + (b_n - a_n)^2}$, we have $d^{(1)}(a_i, b_i) \leq d^{(n)}(\mathbf{a}, \mathbf{b})$, where $a_i = p_i(\mathbf{a})$ and $b_i = p_i(\mathbf{b})$. Therefore $d^{(1)}\big(p_i(\mathbf{a}), p_i(\mathbf{b})\big) \leq d^{(n)}(\mathbf{a}, \mathbf{b}) \leq M$ and so $p_i(A)$ is $d^{(1)}$-bounded.

Solution to Exercise 2.8

Assume that A is bounded. Clearly the result holds for the empty set, so we may assume that A is non-empty and bounded. Then there is an $M > 0$ such that $d(x, y) \leq M$ for all $x, y \in A$.

Fix $a \in A$. Then $d(a, x) \leq M$ for all $x \in A$ and so $d(a, x) < M + \varepsilon$ for all $x \in A$ and for any $\varepsilon > 0$. Therefore $A \subseteq B_d(a, M + \varepsilon)$ for each $\varepsilon > 0$ and so $A \subseteq B_d(a, r)$, where $r = M + \varepsilon$.

Now assume that $A \subseteq B_d(a, r)$ for some $r > 0$. Then for any x and y in A, we have $d(x, y) \leq d(x, a) + d(a, y) \leq r + r = 2r$, so taking $M = 2r$ gives us an $M > 0$ such that $d(x, y) \leq M$ for all $x, y \in A$.

Solution to Exercise 2.9

By Proposition 2.11 it is sufficient to prove that S^3 is a $d^{(4)}$-closed and bounded subset of $\mathbb{R}^4$. Now S^3 is the boundary of $B_{d^{(4)}}(\mathbf{0}, 1)$ and so is $d^{(4)}$-closed. It is clearly bounded and so is sequentially compact.

Solution to Exercise 2.10

Let p_1 and p_2 be the projection functions on $\mathbb{R}^2$. Then p_1 and p_2 are continuous and

$$f(x, y) = p_2((x, y)) \sin(\pi p_1(x, y)) + p_1((x, y)) \cos(\pi p_2(x, y)).$$

It is straightforward to use our list of basic continuous functions from $\mathbb{R}$ to $\mathbb{R}$, together with the combination and Composition Rules for continuous functions from $\mathbb{R}^2$ to $\mathbb{R}$, to verify that

$$(x, y) \mapsto p_2((x, y)) \sin(\pi p_1(x, y))$$

and

$$(x, y) \mapsto p_1((x, y)) \cos(\pi p_2(x, y))$$

are continuous on $\mathbb{R}^2$ and so, by the Sum Rule, f is continuous on $\mathbb{R}^2$.

Hence, by the Restriction Rule, f is continuous on $\{(x, y) : 0 \leq x \leq 1, \text{ and } 0 \leq y \leq 1\}$.

Also, $\{(x, y) : 0 \leq x \leq 1, \text{ and } 0 \leq y \leq 1\}$ is closed and bounded. Hence, by Theorem 2.12, it is sequentially compact.

Thus both conditions of the General Extreme Value Theorem are satisfied, and we conclude that the function f attains a minimum value on $\{(x, y) : 0 \leq x \leq 1, \text{ and } 0 \leq y \leq 1\}$.

Solution to Exercise 3.1

Only if

Let $a \in X$ and suppose that f is (d, e)-continuous at $a \in X$.

Let $\varepsilon > 0$ be given. We suppose, for a contradiction, that for each $\delta > 0$, we can find $x \in B_d(a, \delta)$ for which $e(f(x), f(a)) \geq \varepsilon$. Then, in particular, for each $n \in \mathbb{N}$, we can find $x_n \in B_d(a, \frac{1}{n})$ for which $e(f(x_n), f(a)) \geq \varepsilon$. Thus for this sequence (x_n), we have that it d-converges to a, and yet $e(f(x_n), f(a)) \geq \varepsilon > 0$ for each $n \in \mathbb{N}$, meaning that the sequence $(f(x_n))$ does not e-converge to $f(a)$ – this contradicts the assumption that f is (d, e)-continuous at a.

If

For the opposite implication, let $a \in X$ and suppose that for each $\varepsilon > 0$, there is $\delta > 0$ such that $e(f(x), f(a)) < \varepsilon$, whenever $d(x, a) < \delta$.

Let (x_n) be a sequence in X that d-converges to a. Our aim is to show that $f(x_n)$ e-converges to $f(a)$.

Let $\varepsilon > 0$ be given. Then, by our assumption, we can find $\delta > 0$ for which $e(f(x), f(a)) < \varepsilon$, whenever $d(x, a) < \delta$. However, since $d(x_n, a) \to 0$ as $n \to \infty$, we can find $N \in \mathbb{N}$ so that $d(x_n, a) < \delta$ for $n > N$ and so for $n > N$, $e(f(x_n), f(a)) < \varepsilon$. In other words, for each $\varepsilon > 0$, there is $N \in \mathbb{N}$ so that $e(f(x_n), f(a)) < \varepsilon$ for $n > N$, which means that $f(x_n)$ e-converges to $f(a)$ as $n \to \infty$, and we conclude that f is (d, e)-continuous at a, as required.

Solution to Exercise 3.2

The argument is almost exactly the same as the proof of Theorem 3.3.

We prove the result using proof by contradiction: so suppose that X is sequentially compact and there is $\varepsilon > 0$ such that for each $\delta > 0$ there are $x, y \in X$ for which $d(x, y) < \delta$ and yet $e(f(x), f(y)) \geq \varepsilon$.

Then, in particular, for each $n \in \mathbb{N}$, we can find $x_n, y_n \in X$ for which $d(x_n, y_n) < \frac{1}{n}$ and yet $e(f(x_n), f(y_n)) \geq \varepsilon$.

Now consider the sequence (x_n). Since each term lies in the d-sequentially compact space X, there is a d-convergent subsequence $(x_{n_i})_{i \in \mathbb{N}}$, with limit x, say, that lies in X. Since f is (d, e)-continuous, we know that $f(x_{n_i}) \xrightarrow{d} f(x)$ as $n \to \infty$.

Let us look at the corresponding subsequence (y_{n_i}) of the sequence (y_n). Then we have $d(x_{n_i}, y_{n_i}) < \frac{1}{n_i} \leq \frac{1}{i}$ for each $i \in \mathbb{N}$ and so, since the Triangle Inequality for d implies that

$$d(x, y_{n_i}) \leq d(x, x_{n_i}) + d(x_{n_i}, y_{n_i}) \leq d(x, x_{n_i}) + \frac{1}{i},$$

we deduce, by the Squeeze Rule, that $y_{n_i} \xrightarrow{d} x$ as $i \to \infty$. Hence, since f is continuous, $f(y_{n_i}) \xrightarrow{e} f(x)$ as $i \to \infty$.

However, for each $i \in \mathbb{N}$, the definition of x_{n_i} and y_{n_i} together with the Triangle Inequality for e implies that

$$\begin{aligned} 0 < \varepsilon &\leq e(f(x_{n_i}), f(y_{n_i})) \\ &\leq e(f(x_{n_i}), f(x)) + e(f(x), f(y_{n_i})) \to 0 \text{ as } i \to \infty. \end{aligned}$$

But this is impossible, and we conclude that the theorem holds.

Solution to Exercise 3.3

Suppose that $f\colon X \to Y$ is uniformly continuous. We must show that for each $a \in X$, f is continuous at a, so let $a \in X$ be given.

We use the ε-δ characterisation of continuity given in Theorem 3.2: f is (d, e)-continuous at $a \in X$ if, and only if, for each $\varepsilon > 0$, there is $\delta > 0$ so that $e(f(x), f(a)) < \varepsilon$, whenever $d(x, a) < \delta$.

Fix $\varepsilon > 0$, then since f is uniformly continuous, there is $\delta > 0$ so that for all $x, y \in X$, $d(x, y) < \delta$ implies that $e(f(x), f(y)) < \varepsilon$. In particular, by taking $y = a$, we have the required result.

Solution to Exercise 3.4

The function $f\colon \mathbb{R} \to \mathbb{R}$ given by $f(x) = x^2$ is an example of such a function. It is certainly continuous and for any $x, y \in \mathbb{R}$, we calculate

$$|f(x) - f(y)| = |x - y| \times |x + y|.$$

So if $\varepsilon > 0$ is given, then for any $\delta > 0$, taking $y = x + \frac{1}{2}\delta$ gives $|x - y| = \frac{1}{2}\delta < \delta$ and

$$|f(x) - f(y)| = \tfrac{1}{2}\delta(|2x + \tfrac{1}{2}\delta|).$$

Thus, in order for $|f(x) - f(y)| < \varepsilon$, the choice of δ depends on both x and ε; no single value of $\delta > 0$ works for all possible choices of x. In particular, given $\varepsilon > 0$ and $\delta > 0$, take $x = \frac{\varepsilon}{\delta}$ and $y = x + \frac{1}{2}\delta$, then $|x - y| < \delta$ and

$$|f(x) - f(y)| = \tfrac{1}{2}\delta\left(\left|2\frac{\varepsilon}{\delta} + \tfrac{1}{2}\delta\right|\right) \geq \varepsilon + \left(\tfrac{1}{2}\delta\right)^2 > \varepsilon.$$

Hence f is not uniformly continuous.

Solution to Exercise 3.5

For each $x \in \mathbb{R}$ and $f \in A_g$, $|f(x)| \leq |g(x)|\,(< \infty)$, and so by setting $M_x = |g(x)|$ we see that A_g is pointwise bounded.

Solution to Exercise 3.6

For $x \in [0, 1]$ and each $M > 0$, by considering the constant function $f\colon [0, 1] \to \mathbb{R}$ given by $f(x) = M + 1$, we find a function in A for which $|f(x)| = M + 1 > M$. Hence it is not possible to find a pointwise bound for A at x.

Solution to Exercise 3.7

Fix $x \in X$ and consider $f(x)$ for $f \in A$. Since f has Lipschitz constant L, we can use the Triangle Inequality (for the modulus function) to estimate

$$\begin{aligned} |f(x)| &= |(f(x) - f(a)) + f(a)| \\ &\leq Ld(x,a) + |f(a)| \\ &= Ld(x,a) + |b|, \end{aligned}$$

and this bound is independent of the function f. Hence the set A is pointwise bounded, as required.

Solution to Exercise 3.8

Let $x \in X$ and suppose $M \geq 0$. Then by considering $f\colon X \to \mathbb{R}$ given by $f(x) = M + 1$, which is certainly Lipschitz with Lipschitz constant L and hence in A, we see that $|f(x)| = M + 1 > M$, and so no pointwise bound is possible.

Solution to Exercise 3.9

If A were not pointwise bounded, then there would be $x \in [0,1]$ such that $\sup\{|f(x)| : f \in A\} = \infty$. But this means that we can find a sequence of functions (f_n) for which for all $l \in \mathbb{R}$, there is $N \in \mathbb{N}$ with $|f_n(x) - l| > 1$ for $n > N$, meaning that there could not possibly be a subsequence of (f_n) that converges to a function from $\mathbb{R}$ to $\mathbb{R}$, let alone a continuous one, and this contradicts the sequential compactness of A.

Solution to Exercise 3.10

Suppose that $B \subseteq A$. Then for each $\varepsilon > 0$, there is $\delta > 0$ so that for each $f \in A$, if $x, y \in X$ with $d(x,y) < \delta$, then $e(f(x), f(y)) < \varepsilon$. In particular, it is certainly the case that, since $B \subseteq A$, for each $f \in B$, if $x, y \in X$ with $d(x,y) < \delta$, then $e(f(x), f(y)) < \varepsilon$. Hence B is an equicontinuous set of functions, as required.

Solution to Exercise 3.11

Let $f \in A$.

Since A is equicontinuous, we know that for each $\varepsilon > 0$, there is $\delta > 0$ so that for $g \in A$, if $x, y \in X$ with $d(x,y) < \delta$, then $e(g(x), g(y)) < \varepsilon$. In particular, taking $g = f$, we have if $x, y \in X$ with $d(x,y) < \delta$, then $e(f(x), f(y)) < \varepsilon$. But this is exactly the definition of uniform continuity, and so f is uniformly continuous on X.

Finally, from Lemma 3.6 we conclude that f is continuous on X.

Solution to Exercise 3.12

(a) To show that A is sequentially compact, Theorem 3.12 tells us that it is sufficient to show that:

1. A is closed
2. A is a pointwise bounded family of functions
3. A is an equicontinuous family of functions.

A is closed

Let (f_n) be a sequence of functions in A that converges to $f \in C[0,1]$ for $d_{\max}$. We need to show that $f \in A$.

Since $f_n(0) = 0$ for each $n \in \mathbb{N}$,

$$|f(0)| = |f(0) - f_n(0)| \leq d_{\max}(f_n, f).$$

But $d_{\max}(f_n, f)$ is a null sequence, and so $f(0) = 0$.

For $x, y \in [0,1]$, the Triangle Inequality (for modulus) and the fact that $f_n \in A$ implies that, for each $n \in \mathbb{N}$,

$$\begin{aligned} |f(x) - f(y)| &\leq |f(x) - f_n(x)| + |f_n(x) - f_n(y)| + |f_n(y) - f(y)| \\ &\leq |f_n(x) - f_n(y)| + 2d_{\max}(f_n, f) \\ &\leq \sqrt{|x-y|} + 2d_{\max}(f_n, f). \end{aligned}$$

Since $d_{\max}(f_n, f)$ is a null sequence, it follows that

$$|f(x) - f(y)| \leq \sqrt{|x-y|} \quad \text{for each } x, y \in [0,1].$$

Hence $f \in A$ and so A is closed.

A is a pointwise bounded family of functions

Let $x \in [0,1]$. Then for each $f \in A$,

$$|f(x)| = |f(x) - f(0)| \leq \sqrt{|x-0|} = \sqrt{x} \leq 1.$$

This bound does not depend on the choice of f, and so A is a pointwise bounded family of functions.

A is an equicontinuous family of functions

Let $\varepsilon > 0$. If $x, y \in [0,1]$ and $|x - y| < \delta$ then, for each $f \in A$,

$$|f(x) - f(y)| \leq \sqrt{|x-y|} < \sqrt{\delta}.$$

In particular, if $\delta = \varepsilon^2$, then for each $f \in A$ and each $x, y \in [0,1]$, $|f(x) - f(y)| < \varepsilon$ whenever $x, y \in [0,1]$ satisfies $|x - y| < \delta$. Thus A is an equicontinuous family of functions.

The hypotheses of Theorem 3.12 are satisfied and so A is $d_{\max}$-sequentially compact.

(b) The result of Book D, Chapter 15, Exercise 5.3 tells us that F is $(d_{\max}, d^{(1)})$-continuous and so, since F is a continuous function defined on a sequentially compact set, the General Extreme Value Theorem (Theorem 2.13) implies that F is bounded and attains its bounds on A. In particular, F attains a minimum on A.

Solution to Exercise 4.1

Let $p \in (a, b)$. Then $a < p < b$. Let $r = \min\{|p - a|, |b - p|\}$, then $B_d(p, r) \subseteq (a, b)$.

Solution to Exercise 4.2

One such subcollection is

$$\{(\tfrac{1}{10}(i-1), \tfrac{1}{10}(i+1)) : i = 0, 1, \ldots, 10\}.$$

Solution to Exercise 4.3

(a) If $0 < x \leq 1$, then there is a rational number $q > x$ and so $x \in (0, q)$. Also, $0 \in (-\frac{1}{100}, \frac{1}{100})$.

Hence

$$[0, 1] \subseteq \bigcup_{U \in \mathcal{S}_2} U.$$

Each interval $(0, q)$ is open, and so is the interval $(-\frac{1}{100}, \frac{1}{100})$. Thus $\mathcal{S}_2$ is an open cover of $[0, 1]$.

(b) One such subcollection is the pair of sets

$$\{(-\tfrac{1}{100}, \tfrac{1}{100}), (0, q)\},$$

where q is any fixed rational number strictly larger than 1.

Solution to Exercise 4.4

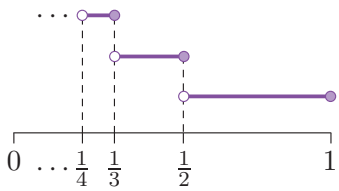

One such collection is

$$\{\{0\}, (\tfrac{1}{2}, 1], (\tfrac{1}{3}, \tfrac{1}{2}], (\tfrac{1}{4}, \tfrac{1}{3}], \ldots, (\tfrac{1}{n+1}, \tfrac{1}{n}], \ldots\}.$$

If we omit even one set, the remaining collection is not a cover of $[0, 1]$. Indeed, whatever set we omit uncovers all the points of that set.

Another possibility is given by

$$\{\{x\} : x \in [0, 1]\}.$$

This is an uncountable cover of $[0, 1]$. If we omit just a single set $\{a\}$ in this cover, then the point $a \in [0, 1]$ is no longer covered.

Solution to Exercise 4.5

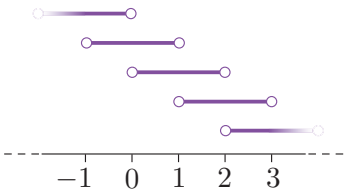

There are many open covers with this property. One such is

$$\mathcal{S} = \{(-1, 1), (0, 2), \ldots, (n-2, n), \ldots\}.$$

If $x \in [0, \infty)$, then x belongs to an interval of the form $(n-2, n)$, for some $n \in \mathbb{N}$. So $\mathcal{S}$ is a cover of $[0, \infty)$. Also, the intervals in $\mathcal{S}$ are all open. So $\mathcal{S}$ is an open cover of $[0, \infty)$. But if

$$\{(n_1 - 2, n_1), \ldots, (n_k - 2, n_k)\}$$

is a finite subcollection of $\mathcal{S}$, then it does not cover any point x for which

$$x \geq \max\{n_1, n_2, \ldots, n_k\}.$$

Solution to Exercise 4.6

Let $\mathcal{S}$ be an arbitrary d-open cover of $\emptyset$. We must find a finite subcover of $\emptyset$ by sets from $\mathcal{S}$. The empty cover $\{\ \}$ is a finite subcover of $\mathcal{S}$ that covers $\emptyset$. We conclude that $\emptyset$ is compact.

Solution to Exercise 4.7

Let (X, d) be a metric space, and let A be a finite subset of X. If $A = \emptyset$, then A is compact by the result of Exercise 4.6. So suppose instead that $A = \{a_1, a_2, \ldots, a_n\} \neq \emptyset$. If $\mathcal{S}$ is an arbitrary d-open cover of A, then, for each $i = 1, 2, \ldots, n$, there is a d-open set $U_i \in \mathcal{S}$ such that $a_i \in U_i$. Hence $\mathcal{R} = \{U_i : i = 1, 2, \ldots, n\}$ is a finite subcover of A from $\mathcal{S}$. We conclude that A is compact.

Solution to Exercise 4.8

We must find an open cover of $\mathbb{R}$ that does not have a finite subcover.

An example of such a cover is given by

$$\mathcal{S} = \{(-n, n) \colon n \in \mathbb{N}\};$$

for if $\mathcal{R}$ is a finite subset of $\mathcal{S}$, then

$$\bigcup_{U \in \mathcal{R}} U = (-N, N)$$

for some $N \in \mathbb{N}$ and does not cover $\mathbb{R}$.

Solution to Exercise 5.1

Let (X, d) be a sequentially compact metric space. Lemma 5.2 implies that there is a finite ε-net $\{x_1, x_2, \ldots, x_N\}$, for some $\varepsilon > 0$. Let

$$M = \max\{d(x_i, x_j) \colon i = 1, 2, \ldots, N \text{ and } j = 1, 2, \ldots, N\}.$$

Now, by the definition of an ε-net, if $x, y \in X$ there are x_i with $d(x, x_i) < \varepsilon$, and x_j with $d(y, x_j) < \varepsilon$, and so

$$d(x, y) \leq d(x, x_i) + d(x_i, x_j) + d(x_j, y) < 2\varepsilon + M.$$

Hence X is bounded.

Solution to Exercise 5.2

(M1) Let $a, b \in X$. Then $d_{\min}(a, b) = \min\{d(a, b), 1\} \geq 0$, since d is a metric. Also since $d(a, a) = 0$ for each $a \in X$, we have $d_{\min}(a, a) = 0$ for each $x \in X$. And if $d_{\min}(a, b) = 0$, then $d(a, b) = 0$ and so $a = b$.

Thus $d_{\min}$ satisfies (M1).

(M2) Let $a, b \in X$. If $d(a, b) \leq 1$ then $d(b, a) \leq 1$ and $d_{\min}(a, b) = d_{\min}(b, a) = d(a, b)$. Also if $d(a, b) > 1$ then $d(b, a) > 1$ and we have $d_{\min}(a, b) = d_{\min}(b, a) = 1$.

Thus, for each $a, b \in X$, (M2) holds.

(M3) Let $a, b, c \in X$. We need to show that $d_{\min}(a, c) \leq d_{\min}(a, b) + d_{\min}(b, c)$. Note that if either $d(a, b)$ or $d(b, c)$ is greater than 1 then $d_{\min}(a, b) + d_{\min}(b, c) \geq 1 \geq d_{\min}(a, c)$ (as $d_{\min}(x, y) \leq 1$ for any x and y). Therefore we may assume that both $d(a, b)$ and $d(b, c)$ are strictly less than 1. In this case, $d_{\min}(a, b) + d_{\min}(b, c) = d(a, b) + d(b, c) \geq d(a, c) \geq d_{\min}(a, c)$ and so (M3) holds.

Solution to Exercise 5.3

First assume that (a_n) is d-convergent to a limit, say $L \in X$. Then, by definition, $(d(a_n, L))$ is a (real) null sequence. To show that (a_n) is $d_{\min}$-convergent, we need to prove that $(d_{\min}(a_n, L))$ is a null sequence. However, $d_{\min}(a, b) \leq d(a, b)$ for any $a, b \in X$ and so, by the Squeeze Rule, (Book D, Chapter 13, Lemma 1.15) the sequence (a_n) is $d_{\min}$-convergent to L.

Now assume that (a_n) is $d_{\min}$-convergent to a limit, say L. Then $(d_{\min}(a_n, L))$ is a null sequence, and so for each $\varepsilon > 0$, there is an $N \in \mathbb{N}$ such that $d_{\min}(a_n, L) < \varepsilon$ whenever $n > N$. We need to prove that $(d(a_n, L))$ is a (real) null sequence. In other words, we need to show that for each $\varepsilon > 0$, there is an $N \in \mathbb{N}$ such that $d(a_n, L) < \varepsilon$ whenever $n > N$.

Let $\varepsilon > 0$ be given and set $\varepsilon' = \min\{\varepsilon, \frac{1}{2}\}$. Then, since (a_n) is $d_{\min}$-convergent, we can find an $N \in \mathbb{N}$ such that $d_{\min}(a_n, L) < \varepsilon'$ whenever $n > N$. Since $\varepsilon' < 1$, we have that $d_{\min}(a_n, L) < 1$ whenever $n > N$, and so we must have $d_{\min}(a_n, L) = d(a_n, L)$ whenever $n > N$. Therefore we have $d(a_n, L) < \varepsilon'$ whenever $n > N$. Hence $(d(a_n, L))$ is a null sequence and (a_n) is d-convergent with limit L.

Solution to Exercise 5.4

We need to show that there is $M > 0$ such that $d_{\min}(x, y) \leq M$ for all $x, y \in A$. Taking $M = 2$ gives the result. (Any number greater than 1 will suffice.)

Solution to Exercise 5.5

We must show that for each $\varepsilon > 0$, there is a finite ε-net. Let $\varepsilon > 0$. If $\varepsilon > (b - a)$, then $\{a\}$ is an ε-net for $[a, b]$. If $0 < \varepsilon \leq (a - b)$, then, for example,

$$\{a + \tfrac{1}{2}n\varepsilon : n \in \mathbb{N} \text{ and } 0 \leq n \leq \tfrac{2(b-a)}{\varepsilon}\}$$

is a finite ε-net for $[a, b]$ with the Euclidean metric.

Solution to Exercise 5.6

Let (X, d) be a compact metric space, and let $\varepsilon > 0$. We seek a finite ε-net for X. The collection

$$\{B_d(x, \varepsilon) : x \in X\}$$

is an open cover of X. Since X is compact, the cover has a finite subcover

$$\{B_d(x_1, \varepsilon), B_d(x_2, \varepsilon), \ldots, B_d(x_n, \varepsilon)\}.$$

So $\{x_1, x_2, \ldots, x_n\}$ is a finite ε-net. Since $\varepsilon > 0$ is arbitrary, X is totally bounded.

Solution to Exercise 6.1

Let $\mathcal{S}$ be an open cover of $\bigcup_{m=1}^{n} K_m$. Then $\mathcal{S}$ is an open cover of K_m for $m = 1, 2, \ldots, n$ and so for each $m = 1, 2, \ldots, n$, we can find a finite subcover of $\mathcal{S}$, $\mathcal{S}_m$ say, such that

$$K_m \subseteq \bigcup_{U \in \mathcal{S}_m} U.$$

But then $\{U : U \in \mathcal{S}_m \text{ for some } m = 1, 2, \ldots, n\}$ is a finite collection of open sets from $\mathcal{S}$ that covers $\bigcup_{m=1}^{n} K_m$. Hence $\mathcal{S}$ has a finite subcover of $\bigcup_{m=1}^{n} K_m$ and the result follows.

Chapter 23

Completeness

Introduction

In the previous two chapters, we have studied connectedness and compactness: in this chapter we discuss a third 'c', known as *completeness*. It comes from an issue with the definition of convergent sequences: recall that a sequence (x_n) in a metric space (X, d) is convergent if there is a point $x \in X$ for which $(d(x_n, x))$ is a null sequence, and the point x is the limit of the sequence. The problem is the following: in order to show that a given sequence is convergent, we need to identify the limit point x, and this can be difficult. So the question we seek to answer in this chapter is:

> Are there metric spaces for which we can identify convergent sequences without first having to identify their limit?

Our answer to this question turns out to be surprisingly useful: the spaces that we identify include many common metric spaces and the property that they have enables us to prove a powerful theorem, the Contraction Mapping Theorem, which is important in the solution of many problems of analysis, some of which we will consider. Indeed, in Chapter 24, *Fractals*, you will see that the Contraction Mapping Theorem plays an important role in the investigation of fractals.

Chapter guide

We first explain what we mean by a sequence that 'appears to converge'. We then identify those metric spaces for which such sequences *do* converge – we call these *complete* metric spaces.

We then discuss the Contraction Mapping Theorem – a powerful theorem that gives sufficient conditions to ensure that a continuous mapping of a certain type on a complete metric space must have a fixed point. The Contraction Mapping Theorem is a surprisingly useful result and we give several examples of its applications.

In the last section of the chapter, we show how to embed any metric space into a complete metric space.

1 Completeness

In analysis, we are frequently given a sequence in a metric space, and we wish to know whether it is convergent.

In Book D, Chapter 14, we defined a convergent sequence as follows:

Definition 1.1 *Definition 3.2 of Chapter 14*

Let (X, d) be a metric space. A sequence (a_k) in X **d-converges** to $a \in X$ if the sequence of real numbers $(d(a_k, a))$ is a null sequence.

When using this definition to show that a given sequence is convergent, we must first identify a possible limit of the sequence and then show that the sequence converges to it. In many instances, however, it is difficult to identify a possible limit. For example, consider the sequence of real numbers (a_n) given by

$$a_n = 1 + \frac{1}{2^2} + \frac{1}{3^2} + \cdots + \frac{1}{n^2}.$$

See Book D, Chapter 13, Theorem 1.22.

The limit of this sequence is not obvious. Fortunately, we have the Monotone Convergence Theorem available, which states that a non-decreasing sequence of real numbers that is bounded above is convergent. This sequence is non-decreasing (since $a_{n+1} - a_n = \frac{1}{(n+1)^2} > 0$ for $n \in \mathbb{N}$) and is bounded above by 2 (since $a_n \leq 1 + \int_1^n \frac{1}{x^2}\,dx$). So, by the Monotone Convergence Theorem, we know that (a_n) converges even though we have not identified its limit.

The limit of this sequence is in fact $\frac{\pi^2}{6}$ – this was first calculated in the eighteenth century by Leonhard Euler.

Now consider the sequence (b_n) where

$$b_n = -1 + \frac{1}{2} - \frac{1}{3} + \cdots + \frac{(-1)^n}{n}.$$

In this case, $b_{n+1} - b_n = \frac{(-1)^{n+1}}{n+1}$, and so the sequence is neither monotonic increasing nor decreasing. However, it is possible to show the sequence is convergent.

Exercise 1.1

Let $\mathbb{R}$ have the Euclidean metric and define the sequence (b_n) by $b_n = -1 + \frac{1}{2} - \frac{1}{3} + \cdots + \frac{(-1)^n}{n}$.

(a) Show that the odd and even subsequences of (b_n), $(b_{2n-1})_{n\in\mathbb{N}}$ and $(b_{2n})_{n\in\mathbb{N}}$, are bounded monotonic sequences.

(b) Show that the sequence (d_n) given by $d_n = b_{2n-1} - b_{2n}$ is null.

(c) Hence show that the sequence (b_n) is convergent.

We have now seen two examples of convergent sequences in $\mathbb{R}$ for which we were able to prove convergence without needing to identify the limit. But, as we saw for the second sequence, proving convergence was fairly involved. What we would like to do is to find some property of convergent sequences that we should test for instead.

By considering a sequence such as $(1/n)$, which converges in $\mathbb{R}$ with the Euclidean metric but diverges in $(0, 1]$ with the (induced) Euclidean metric, we see that there will be an interaction between the property that

we identify and the metric space that we are working with. As a result, we identify three goals that are closely intertwined:

1. identify a property, P say, that all convergent sequences (in any metric space) possess
2. determining whether a given sequence has property P should not require identification of a potential limit for the sequence
3. for a large class of metric spaces, the set of sequences with property P is the same as the set of convergent sequences – that is, for these metric spaces, a sequence should be convergent if, and only if, it has property P. Ideally the class should include metric spaces that are of use to us, such as Euclidean spaces and $C[0,1]$ with the max metric.

In both of the example sequences (a_n) and (b_n) that we considered above, we made extensive use of the Monotone Convergence Theorem, and it would be reasonable to consider this as our starting point. Unfortunately, the Monotone Convergence Theorem (as its name suggests) makes extensive use of the fact that $\mathbb{R}$ can be ordered – by which we mean two real numbers can be compared: $x < y$, $x = y$ or $x > y$. This is a very strong property and not appropriate for generalisation here. (For example, it is hard to think of a natural way of ordering points in the plane that is compatible with our usual notion of distance.)

The property that we work with instead is based around a different idea: the distances between the terms of a convergent sequence. We discuss it in the next subsection.

1.1 Cauchy sequences

Suppose that (a_n) is a convergent sequence (in a metric space (X,d)) and let $a \in X$ be its limit. Then we can use the Triangle Inequality to deduce that for $n, m \in \mathbb{N}$

$$d(a_n, a_m) \leq d(a_n, a) + d(a, a_m) = d(a_n, a) + d(a_m, a).$$

But $(d(a_n, a))$ and $(d(a_m, a))$ are (real) null sequences, and so by choosing n, m large enough, we can make $d(a_n, a_m)$ arbitrarily small. The following lemma states this observation more precisely.

Lemma 1.2

Let (X,d) be a metric space and let (a_n) be a convergent sequence in X. Then, for each $\varepsilon > 0$, there is an $N \in \mathbb{N}$ such that

$$d(a_n, a_m) < \varepsilon \text{ for each } n, m > N.$$

Proof Let (a_n) be a convergent sequence in X. Then there is an $a \in X$ such that

$$a_n \xrightarrow{d} a \text{ as } n \to \infty.$$

Let $\varepsilon > 0$. Then there is an $N \in \mathbb{N}$ such that, for each $n > N$,

$$d(a_n, a) < \tfrac{1}{2}\varepsilon.$$

We deduce from the Triangle Inequality that, for all $n, m > N$,

$$d(a_n, a_m) \leq d(a_n, a) + d(a, a_m) < \tfrac{1}{2}\varepsilon + \tfrac{1}{2}\varepsilon = \varepsilon,$$

as required. ■

The property of convergent sequences given in Lemma 1.2 is the one we seek.

Augustin-Louis Cauchy (1789–1857) was a prolific French mathematician whose name occurs in many areas of analysis. His book *Cours d'analyse* was one of the first to place the foundations of calculus on a rigorous basis.

Definition 1.3 *Cauchy sequence*

Let (X, d) be a metric space and let (a_n) be a sequence in X.

Then (a_n) is a **d-Cauchy sequence** if, for each $\varepsilon > 0$, there is an $N \in \mathbb{N}$ such that

$$d(a_n, a_m) < \varepsilon \text{ for each } n, m > N.$$

Remarks

1. When it is clear from the context, we suppress the d and just refer to a sequence (a_n) as being a Cauchy sequence.
2. Lemma 1.2 states that any convergent sequence in a metric space is a Cauchy sequence.
3. Thus, since they are convergent for the Euclidean metric, the sequences (a_n) given by $a_n = 1 + \frac{1}{2^2} + \frac{1}{3^2} + \cdots + \frac{1}{n^2}$ and (b_n) given by $b_n = -1 + \frac{1}{2} - \frac{1}{3} + \cdots + \frac{(-1)^n}{n}$ are $d^{(1)}$-Cauchy.
4. This definition does not require us to identify a limit of the sequence, or even to know whether the sequence *has* a limit or not.
5. We may assume that $m > n$ in the definition, since if $n = m$, then $d(a_n, a_m) = 0$ which is certainly smaller than ε, and if $m < n$, we can interchange n and m since $d(a_n, a_m) = d(a_m, a_n)$.
6. To show that a given sequence is *not* a Cauchy sequence, it is enough to find just *one* $\varepsilon > 0$ such that, for each $N \in \mathbb{N}$, there are n and m both larger than N with $d(a_n, a_m) \geq \varepsilon$.
7. If we wish to emphasise the metric, we say that (a_n) is a Cauchy sequence *for* d or *with respect to* d, or that (a_n) is a d-Cauchy sequence.

Exercise 1.2

Determine whether each of the following sequences is a d-Cauchy sequence in $\mathbb{R}$, where d denotes the Euclidean metric.

(a) $(1/n)$ (b) (n)

Exercise 1.3

Let (X, d) be a metric space and (a_n) an eventually constant sequence in X. Show that (a_n) is a d-Cauchy sequence.

We now derive some basic properties of Cauchy sequences.

Lemma 1.4

Let (X, d) be a metric space and let (a_n) be a d-Cauchy sequence in X. Then (a_n) is a bounded sequence in X – that is, there are an $a \in X$ and an $M > 0$ such that

$$\{a_n : n \in \mathbb{N}\} \subseteq B_d(a, M).$$

Proof Since (a_n) is a Cauchy sequence, there is an $N \in \mathbb{N}$ such that

$$d(a_n, a_m) < 1 \text{ for each } n, m > N.$$

Let $a = a_{N+1}$; then

$$d(a_n, a) < 1 \text{ whenever } n > N.$$

So, for each $n \in \mathbb{N}$,

$$d(a_n, a) \leq \max\{1, d(a_1, a), d(a_2, a), d(a_3, a), \ldots, d(a_N, a)\}.$$

This implies that

$$\{a_n : n \in \mathbb{N}\} \subseteq B_d(a, M),$$

where $M = 1 + \max\{1, d(a_1, a), d(a_2, a), d(a_3, a), \ldots, d(a_N, a)\}$. ■

Corollary 1.5

Each $d^{(1)}$-Cauchy sequence in $\mathbb{R}$ is bounded.

Proof Let (a_n) be a $d^{(1)}$-Cauchy sequence in $\mathbb{R}$. By Lemma 1.4, there are an $a \in \mathbb{R}$ and an $M > 0$ such that

$$|a_n - a| < M \text{ for each } n \in \mathbb{N}.$$

By the Triangle Inequality,

$$|a_n| = |a_n - a + a| \leq |a_n - a| + |a| < M + |a| \text{ for each } n \in \mathbb{N}.$$

So (a_n) is bounded. ■

We now show that if a subsequence of a Cauchy sequence converges, then the whole sequence must converge. This is not true for sequences in general – for example, the sequence $(1, -1, 1, -1, \ldots)$ does not converge (for the Euclidean metric) but has the convergent subsequences $(1, 1, 1, \ldots)$ and $(-1, -1, -1, \ldots)$.

Lemma 1.6

Let (X, d) be a metric space and let (a_n) be a Cauchy sequence in X. If (a_n) has a convergent subsequence with limit a in X, then (a_n) is convergent with limit a.

Proof Let (a_{n_k}) be a subsequence of (a_n) that d-converges to $a \in X$. We prove that (a_n) also d-converges to a.

Let $\varepsilon > 0$ be given. Since $a_{n_k} \to a$ as $k \to \infty$, there is a $K \in \mathbb{N}$ such that for each $k > K$,

$$d(a_{n_k}, a) < \tfrac{1}{2}\varepsilon.$$

Since (a_n) is a d-Cauchy sequence, there is an $N \in \mathbb{N}$ such that

$$d(a_n, a_m) < \tfrac{1}{2}\varepsilon \text{ for each } m, n > N.$$

Let $k > K$ be chosen so that $n_k > N$; then

$$d(a_n, a_{n_k}) < \tfrac{1}{2}\varepsilon \text{ for each } n > N.$$

Moreover, since $k > K$,

$$d(a_{n_k}, a) < \tfrac{1}{2}\varepsilon.$$

Hence by the Triangle Inequality for d, if $n > N$, then

$$d(a_n, a) \leq d(a_n, a_{n_k}) + d(a_{n_k}, a) < \tfrac{1}{2}\varepsilon + \tfrac{1}{2}\varepsilon = \varepsilon \text{ for each } n > N.$$

Since $\varepsilon > 0$ is arbitrary, (a_n) converges to a in (X, d). ■

Remark

Thus for the sequence (b_n) given by $b_n = -1 + \frac{1}{2} - \frac{1}{3} + \cdots + \frac{(-1)^n}{n}$, an alternative way to prove that it is convergent for the Euclidean metric is to show that (b_n) is a Cauchy sequence and then show that it has a convergent subsequence.

Exercise 1.4

Let (c_n) be the sequence of real numbers given by $c_n = \sum_{i=1}^{n} \frac{(-1)^i}{2^i}$.

(a) Show that (c_n) is a Cauchy sequence for the Euclidean metric.

(b) By considering the subsequence of even terms, show that (c_n) is a convergent sequence.

1.2 Cauchy sequences in $\mathbb{R}$

We have seen that in any metric space, each convergent sequence is a Cauchy sequence (Lemma 1.2). We now show that in $\mathbb{R}$ (with its usual metric) the Cauchy sequences are exactly the convergent sequences, thus providing further evidence that the Cauchy property is the one we seek. Since we already know from Lemma 1.2 that every convergent sequence in $\mathbb{R}$ is a Cauchy sequence, we need only now show that each Cauchy sequence is convergent in $\mathbb{R}$ for the Euclidean metric.

To do this, we recall from Corollary 1.5 that a Cauchy sequence in $\mathbb{R}$ is bounded, and then construct a monotonic subsequence of the Cauchy sequence. By the Monotone Convergence Theorem, this subsequence is convergent, and so by Lemma 1.6 the original Cauchy sequence is convergent. We need the following result that we proved in Chapter 22.

See Book D, Chapter 13, Theorem 1.22.

Lemma 1.7 *Lemma 1.1 of Chapter 22*

Let (a_n) be a sequence of real numbers. Then (a_n) contains a monotonic subsequence.

We are now in a position to show that every Cauchy sequence in $\mathbb{R}$ is convergent.

Theorem 1.8

Each $d^{(1)}$-Cauchy sequence in $\mathbb{R}$ is $d^{(1)}$-convergent.

Proof Let (a_n) be a Cauchy sequence of real numbers.

By Corollary 1.5, there is a positive real number M such that

$$|a_n| < M \text{ for each } n \in \mathbb{N}.$$

By Lemma 1.7, there exists a monotonic subsequence (a_{n_k}) of (a_n). Since (a_n) is a bounded sequence, the monotonic subsequence is also bounded. Hence, by the Monotone Convergence Theorem, the subsequence is convergent. Thus, by Lemma 1.6, (a_n) is convergent. ■

Combining this result with our earlier observation that convergent sequences in metric spaces are Cauchy sequences, we obtain the following result.

Corollary 1.9

A sequence in $\mathbb{R}$ is $d^{(1)}$-convergent if, and only if, it is a $d^{(1)}$-Cauchy sequence.

We have shown that the $d^{(1)}$-Cauchy sequences in $\mathbb{R}$ are precisely the $d^{(1)}$-convergent sequences. This enables us to determine whether a sequence in $\mathbb{R}$ is convergent without needing to identify the limit of the sequence – exactly what we wanted do. Moreover, we have shown that in any metric space, convergent sequences are Cauchy and so we have a candidate for our property P: being Cauchy. It remains to show that there are useful metric spaces for which being Cauchy is the same as being convergent.

In the next section, we show that many of our standard examples of metric spaces have exactly this property and we develop some methods for identifying metric spaces for which Cauchy sequences correspond to the convergent sequences.

2 Complete metric spaces

We have just seen that $\mathbb{R}$ with the usual Euclidean metric is an example of a metric space for which a sequence is convergent if, and only if, it is a Cauchy sequence. Our aim in this section is to develop techniques that enable us to identify other examples of spaces for which the Cauchy sequences are exactly the convergent sequences.

We start by giving a name to those metric spaces for which each Cauchy sequence is convergent.

Definition 2.1 *Complete metric space*

A metric space (X, d) is a **complete metric space** if each Cauchy sequence in (X, d) is convergent; otherwise it is an **incomplete metric space**.

Remarks

1. Note that the definition of complete requires convergence to a point of X. Exercise 2.1 emphasises this subtlety.
2. Theorem 1.8 tells us that $(\mathbb{R}, d^{(1)})$ is a complete metric space.

Exercise 2.1

Show that $(0, 1]$ with the induced Euclidean metric is incomplete.

Exercise 2.2

Let X be a set and let d_0 denote the discrete metric for X. Show that (X, d_0) is a complete metric space.

Thus we have at least two examples of complete metric spaces: $(\mathbb{R}, d^{(1)})$ and (X, d_0), for any set X. We now develop criteria that guarantee that a given space is complete and identify some other complete metric spaces.

2.1 Criteria for completeness

We start by investigating when subspaces of a metric space are complete. We have already seen that, although $(\mathbb{R}, d^{(1)})$ is complete, the subspace $(0, 1]$ with the induced Euclidean metric is incomplete since there are Cauchy sequences in $(0, 1]$ that do not converge within $(0, 1]$.

The sequence $(\frac{1}{n})$ is such a sequence.

Our next result shows that $(A, d_A^{(1)})$ is complete if A is any *closed* subset of $\mathbb{R}$. The proof relies on the following characterisation of the closure of a set in terms of closure points:

Worked Exercise 3.7 and Theorem 3.8 of Chapter 16

Let (X, d) be a metric space and $A \subseteq X$. Then

$$\mathrm{Cl}_{(X,d)}(A) = \{x \in X : x \text{ is a } d\text{-closure point of } A \text{ in } X\},$$

where $x \in X$ is a d-closure point of A if there is a sequence in A that d-converges to x.

The following proposition uses both the induced metric and the original metric, and care is required to ensure that the right metric is used at the right time.

Proposition 2.2

Let (X, d) be a metric space and let $A \subseteq X$. Let d_A denote the induced metric on A from d.

1. If (A, d_A) is a complete metric space, then A is d-closed.
2. If (X, d) is a complete metric space and A is d-closed, then (A, d_A) is a complete metric space.

Proof

1. *Let (A, d_A) be complete. We show that A is d-closed.*

   Since a set is d-closed if, and only if, it coincides with its d-closure (see Theorem 3.3 of Chapter 16), and since it is always the case that $A \subseteq \mathrm{Cl}_{(X,d)}(A)$, it is enough to show that $\mathrm{Cl}_{(X,d)}(A) \subseteq A$. So let $x \in \mathrm{Cl}_{(X,d)}(A)$, and let (a_n) be a sequence in A that d-converges to x in X. By Lemma 1.2, (a_n) is a d-Cauchy sequence in X, and so it is a d_A-Cauchy sequence in A. But (A, d_A) is complete, and so (a_n) converges to a point in A. Since limits are unique in a metric space, this limit is x. Thus $x \in A$ so $\mathrm{Cl}_{(X,d)}(A) \subseteq A$.

See Book D, Chapter 14, Theorem 3.3.

2. *Let (X, d) be complete and let $A \subseteq X$ be d-closed. We show that (A, d_A) is complete.*

   If (a_n) is a d_A-Cauchy sequence in A, then it is also a d-Cauchy sequence in X. Since X is complete, (a_n) d-converges to some $x \in X$. Since each term of the sequence (a_n) is in A, the d-limit x belongs to $\mathrm{Cl}_{(X,d)}(A)$. Since A is d-closed, $A = \mathrm{Cl}_{(X,d)}(A)$ and so $x \in A$, meaning that (a_n) is d_A-convergent. Thus (A, d_A) is complete.

■

It follows immediately from Proposition 2.2 that, if $A \subseteq \mathbb{R}$, then $(A, d_A^{(1)})$ is a complete metric space if, and only if, A is $d^{(1)}$-closed.

Exercise 2.3

For which of the following sets A is $(A, d_A^{(1)})$ complete?

(a) $A = [0, 1]$

(b) $A = [0, 1)$

(c) $A = (0, 1)$

We now look at the relationship between completeness and compactness. Recall from Theorem 5.12 of Chapter 22 that metric spaces are compact if, and only if, they are sequentially compact. Since compact subsets of a metric space are closed, Proposition 2.2 also implies that compact subsets of complete metric spaces are complete. In fact, we can say more:

See Corollary 5.13 part 1 of Chapter 22.

Theorem 2.3

Each compact metric space is complete.

Proof Let (X, d) be a compact metric space and suppose that (a_n) is a Cauchy sequence in X. Since X is compact, then it is sequentially compact and hence (a_n) has a d-convergent subsequence, and so we conclude from Lemma 1.6 that (a_n) is d-convergent, as required. ■

Exercise 2.4

Let (X, d) be a compact metric space, let (Y, e) be a metric space, and let $f\colon X \to Y$ be a (d, e)-continuous onto function.

(a) Is (Y, e) complete?

(b) What can you deduce if f is not onto?

We now ask you to show that a finite metric space is complete.

Exercise 2.5

Let $X = \{x_1, x_2, \ldots, x_n\}$ be a finite set, and let d be a metric on X. Prove that (X, d) is complete.

This is a particularly simple example of a complete metric space. Our remaining examples of complete metric spaces are more interesting.

2.2 The completeness of $(\mathbb{R}^n, d^{(n)})$

In Theorem 1.8 we showed that $(\mathbb{R}, d^{(1)})$ is a complete metric space. We now show that $(\mathbb{R}^n, d^{(n)})$ is a complete metric space, for any $n \geq 1$.

In Chapter 14, we saw that a sequence $(\mathbf{a}_k)$ in $\mathbb{R}^n$ converges to $\mathbf{a}$ with respect to the Euclidean metric $d^{(n)}$ if, and only if, each component sequence $(a_{j,k})_{k\in\mathbb{N}}$ converges to a_j, for $j = 1, 2, \ldots, n$. Using the same idea, we can show that $(\mathbb{R}^n, d^{(n)})$ is complete – that is, each $d^{(n)}$-Cauchy sequence in $\mathbb{R}^n$ converges to a point in $\mathbb{R}^n$.

See Book D, Chapter 14, Proposition 1.9.

Here $\mathbf{a} = (a_1, a_2, \ldots, a_n)$ and $\mathbf{a}_k = (a_{1,k}, a_{2,k}, \ldots, a_{n,k})$.

We first consider the relationship between a Cauchy sequence in $\mathbb{R}^n$ and the corresponding component sequences.

Lemma 2.4

A sequence $(\mathbf{a}_k)_{k\in\mathbb{N}}$ is a $d^{(n)}$-Cauchy sequence in $\mathbb{R}^n$ if, and only if, for each $j = 1, 2, \ldots, n$ the component sequence $(a_{j,k})_{k\in\mathbb{N}}$ is a $d^{(1)}$-Cauchy sequence in $\mathbb{R}$.

Proof Let $(\mathbf{a}_k)_{k\in\mathbb{N}}$ be a $d^{(n)}$-Cauchy sequence in $\mathbb{R}^n$. Given $\varepsilon > 0$, choose a $K \in \mathbb{N}$ such that

$$d^{(n)}(\mathbf{a}_k, \mathbf{a}_l) < \varepsilon \text{ for each } k, l > K.$$

For each $j \in \{1, 2, \ldots, n\}$,

$$|a_{j,k} - a_{j,l}| \leq d^{(n)}(\mathbf{a}_k, \mathbf{a}_l) < \varepsilon,$$

and so for each $j = 1, 2, \ldots, n$, $(a_{j,k})_{k\in\mathbb{N}}$ is a $d^{(1)}$-Cauchy sequence.

Conversely, suppose that $(a_{j,k})_{k\in\mathbb{N}}$ is a $d^{(1)}$-Cauchy sequence in $\mathbb{R}$ for each $j = 1, 2, \ldots, n$. Given $\varepsilon > 0$ and $j \in \{1, 2, \ldots, n\}$, choose a $K_j \in \mathbb{N}$ such that

$$|a_{j,k} - a_{j,l}| < \tfrac{1}{\sqrt{n}}\varepsilon \text{ for each } k, l > K_j.$$

Then, for $k, l > \max\{K_1, K_2, \ldots, K_n\}$,

$$d^{(n)}(\mathbf{a}_k, \mathbf{a}_l)^2 = \sum_{j=1}^{n} |a_{j,k} - a_{j,l}|^2 < \sum_{j=1}^{n} \left(\tfrac{1}{\sqrt{n}}\varepsilon\right)^2 = \varepsilon^2$$

and so $d^{(n)}(\mathbf{a}_k, \mathbf{a}_l) < \varepsilon$. Thus $(\mathbf{a}_k)_{k\in\mathbb{N}}$ is a $d^{(n)}$-Cauchy sequence in $\mathbb{R}^n$. ■

This lemma allows us to show that $(\mathbb{R}^n, d^{(n)})$ is complete.

Corollary 2.5

The metric space $(\mathbb{R}^n, d^{(n)})$ is complete.

Proof Let $(\mathbf{a}_k)_{k\in\mathbb{N}}$ be a $d^{(n)}$-Cauchy sequence in $\mathbb{R}^n$. By Lemma 2.4, $(a_{j,k})_{k\in\mathbb{N}}$ is a $d^{(1)}$-Cauchy sequence in $\mathbb{R}$, for $j = 1, 2, \ldots, n$, and so, since $\mathbb{R}$ is complete (see Theorem 1.8), it converges to $a_j \in \mathbb{R}$, say. Thus $(\mathbf{a}_k)$ is convergent (and converges to $\mathbf{a} = (a_1, a_2, \ldots, a_n)$). Hence $(\mathbb{R}^n, d^{(n)})$ is complete. ■

See Book D, Chapter 14, Proposition 1.9.

Exercise 2.6

Let $A = \{\mathbf{x} \in \mathbb{R}^n : d^{(n)}(\mathbf{x}, \mathbf{0}) \leq 1\}$. Show that $(A, d_A^{(n)})$ is a complete metric space.

2.3 Cantor space is complete

We now show that the Cantor space (with the Cantor metric) is a complete metric space. First let us recall the definition of Cantor space.

Cantor space (see Definitions 2.1 and 2.2 of Chapter 15)

The Cantor space $\mathbf{C}$ consists of all infinite sequences of zeros and ones:

$$\mathbf{C} = \{(a_n) : a_n \in \{0, 1\} \text{ for each } n \in \mathbb{N}\},$$

and the Cantor metric $d_{\mathbf{C}} : \mathbf{C} \times \mathbf{C} \to \mathbb{R}$ is given by

$$d_{\mathbf{C}}(\mathbf{x}, \mathbf{y}) = \begin{cases} 0, & \text{if } \mathbf{x} = \mathbf{y}, \\ 2^{-n}, & \text{if } \mathbf{x} \text{ and } \mathbf{y} \text{ first differ at their } n\text{th term.} \end{cases}$$

The fact that Cantor space is complete follows from the following observation concerning $d_{\mathbf{C}}$-Cauchy sequences.

Lemma 2.6

Let $(\mathbf{x}_n)$ be a $d_{\mathbf{C}}$-Cauchy sequence in $\mathbf{C}$. Then for each $k \in \mathbb{N}$, there is $N_k \in \mathbb{N}$ such that for $n, m > N_k$,

$$x_{i,n} = x_{i,m} \text{ for } i = 1, \ldots, k.$$

In other words, the initial entries of terms in a Cauchy sequence are eventually the same.

Exercise 2.7

Prove Lemma 2.6.

Theorem 2.7

$(\mathbf{C}, d_{\mathbf{C}})$ is a complete metric space.

Proof Let $(\mathbf{x}_n)$ be a $d_{\mathbf{C}}$-Cauchy sequence in $\mathbf{C}$. We must show that it is a $d_{\mathbf{C}}$-convergent sequence. To do this, we need to identify the limit of the sequence, $\mathbf{x}$ say, and then show that $\mathbf{x}_n \xrightarrow{d_{\mathbf{C}}} \mathbf{x}$ as $n \to \infty$.

Finding a candidate for the limit of $(\mathbf{x}_n)$

For $k \in \mathbb{N}$, let $N_k \in \mathbb{N}$ be as given in Lemma 2.6. We define the candidate limit sequence $\mathbf{x} = (x_1, x_2, x_3, \ldots)$ as follows:

$$x_k = x_{k,N_k+1} \text{ for each } k \in \mathbb{N}.$$

Showing that $\mathbf{x}_n \to \mathbf{x}$

Let $\varepsilon > 0$ and choose $k \in \mathbb{N}$ so that $2^{-k} < \varepsilon$, then for $n > N_k$, we have

$$x_{i,n} = x_i \text{ for each } i = 1, 2, \ldots, k$$

and so

$$d_{\mathbf{C}}(\mathbf{x}_n, \mathbf{x}) \leq 2^{-k} < \varepsilon.$$

Hence $(\mathbf{x}_n)$ converges to $\mathbf{x}$ for $d_{\mathbf{C}}$, as required. ■

Remark

Most of the effort here was concentrated on identifying the limit of the Cauchy sequence. Once a candidate limit had been identified, the verification that it was indeed the limit was straightforward.

2.4 $(C[0,1], d_{\max})$ is complete

We now turn to an important example of a complete metric space that is not a subset of a Euclidean space: $(C[0,1], d_{\max})$.

We have already studied this space in some detail. Recall that $d_{\max}$ is given by

$$d_{\max}(f, g) = \max\{|g(x) - f(x)| : x \in [0,1]\},$$

(Figure 2.1) and measures the maximum difference between the values of two continuous functions on $[0,1]$.

In this subsection, we prove the following result.

Theorem 2.8

The metric space $(C[0,1], d_{\max})$ is complete.

See Book D, Chapter 15, Section 5.

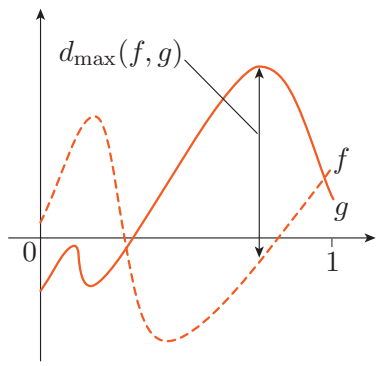

Figure 2.1

The proof of this result is sophisticated, and in order to prove it we use the notion of uniform convergence that we introduced in Section 4 of Chapter 15.

Recall that uniform convergence for a sequence of functions from $[0,1]$ to $\mathbb{R}$ is defined as follows.

Uniform convergence (see Definition 4.4 of Chapter 15)

A sequence (f_n) of functions $f_n\colon [0,1] \to \mathbb{R}$ **converges uniformly** on $[0,1]$ to the function $f\colon [0,1] \to \mathbb{R}$ if

1. the function $f_n - f$ is bounded for each $n \in N$, and
2. the sequence (M_n) defined by
$$M_n = \sup\{|f_n(x) - f(x)| : x \in [0,1]\}$$
is a (real) null sequence.

Proof of Theorem 2.8 Let (f_n) be a $d_{\max}$-Cauchy sequence of functions in $C[0,1]$. The proof consists of three steps:

1. identify a candidate limit function, f, for the sequence (f_n)
2. show that (f_n) converges uniformly to f
3. deduce that f is continuous and that $d_{\max}(f_n, f) \to 0$ as $n \to \infty$.

If we do this, then we can conclude that (f_n) is $d_{\max}$-convergent to f in $C[0,1]$ and so $(C[0,1], d_{\max})$ is complete.

Identify a candidate limit function, f, of the sequence (f_n)

For each $x \in [0,1]$ and for each $n, m \in \mathbb{N}$,

$$|f_n(x) - f_m(x)| = d^{(1)}(f_n(x), f_m(x)) \leq d_{\max}(f_n, f_m).$$

Thus, for each $x \in [0,1]$, $(f_n(x))$ is a $d^{(1)}$-Cauchy sequence in $\mathbb{R}$. By Theorem 1.8, $(f_n(x))$ is $d^{(1)}$-convergent, and we define $f\colon [0,1] \to \mathbb{R}$ by

$$f(x) = \lim_{n\to\infty} f_n(x).$$

This function f is our candidate for the $d_{\max}$-limit of the sequence (f_n).

Show that (f_n) converges uniformly to f

We first show that the function $f_n - f$ is bounded for each $n \in \mathbb{N}$. To do this, we observe that Lemma 1.4 implies that there are an $M > 0$ and a $g \in C[0,1]$ such that

$$f_n \in B_{d_{\max}}(g, M) \text{ for each } n \in \mathbb{N}.$$

Hence for each $x \in [0,1]$ and each $n \in \mathbb{N}$,

$$|f_n(x) - g(x)| \leq d_{\max}(f_n, g) < M.$$

The Triangle Inequality now implies that, for each $x \in [0,1]$ and each $n \in \mathbb{N}$,

$$|f(x) - g(x)| \leq |f(x) - f_n(x)| + |f_n(x) - g(x)| < |f(x) - f_n(x)| + M.$$

For each $x \in [0,1]$, $|f(x) - f_n(x)| \to 0$ as $n \to \infty$, and so

$$|f(x) - g(x)| \leq M.$$

Using the Triangle Inequality again, we find that for each $x \in [0,1]$,

$$|f(x) - f_n(x)| \leq |f(x) - g(x)| + |g(x) - f_n(x)| \leq 2M.$$

Thus the function $f_n - f$ is bounded for each $n \in \mathbb{N}$.

We now show that

$$M_n = \sup\{|f_n(x) - f(x)| : x \in [0,1]\}$$

is a null sequence and so, by the definition of uniform convergence, (f_n) converges uniformly to f.

Let $\varepsilon > 0$. Since (f_n) is a $d_{\max}$-Cauchy sequence, there is an $N \in \mathbb{N}$ such that for each $m, n > N$,

$$d_{\max}(f_n, f_m) < \tfrac{1}{2}\varepsilon.$$

Let $x \in [0,1]$. Since $(f_n(x))$ converges to $f(x)$, there is an $M \in \mathbb{N}$ such that for each $m > M$,

$$|f(x) - f_m(x)| < \tfrac{1}{2}\varepsilon.$$

Note the introduction of the function f_m and the consequent disappearance of the dependence on x in the bound for $|f(x) - f_n(x)|$. This is a *uniformisation argument*: we find an estimate of $|f(x) - f_n(x)|$ that is independent of the choice of X. This is a useful technique.

Let $n > N$ and $m > \max\{M, N\}$. We use the Triangle Inequality to estimate that

$$\begin{aligned} |f(x) - f_n(x)| &\leq |f(x) - f_m(x)| + |f_m(x) - f_n(x)| \\ &\leq |f(x) - f_m(x)| + d_{\max}(f_n, f_m) < \tfrac{1}{2}\varepsilon + \tfrac{1}{2}\varepsilon = \varepsilon. \end{aligned}$$

Since $x \in [0,1]$ is arbitrary, we conclude that for $n > N$,

$$M_n = \sup\{|f(x) - f_n(x)| : x \in [0,1]\} \leq \varepsilon.$$

But $\varepsilon > 0$ is arbitrary and so (M_n) is a null sequence. Thus (f_n) converges uniformly to f.

Deduce that $f \in C[0,1]$ *and that* $d_{\max}(f_n, f) \to 0$ *as* $n \to \infty$

Since (f_n) converges uniformly to f, the Uniform Convergence Theorem (Theorem 5.8 of Chapter 15) implies that f is continuous on $[0,1]$, and so is an element of $C[0,1]$, and $d_{\max}(f_n, f) \to 0$ as $n \to \infty$.

Hence (f_n) is $d_{\max}$-convergent to f in $C[0,1]$, as required. ■

In the next example, we give a metric on $C[0,1]$ for which the resulting space is *not* complete.

Example 2.9 $C[0,1]$ *with the integral metric*

The integral metric on $C[0,1]$ is given by

$$d(f,g) = \int_0^1 |g(x) - f(x)|\, dx.$$

We saw in Book D, Chapter 15, Worked Exercise 5.3 that this defines a metric on $C[0,1]$.

Define a sequence of functions (f_n) in $C[0,1]$ by setting, for $n \in \mathbb{N}$,

$$f_n(x) = \begin{cases} 0, & \text{if } 0 \le x \le \frac{1}{2}\left(1 - \frac{1}{n}\right), \\ 2n\left(x - \frac{1}{2}\left(1 - \frac{1}{n}\right)\right), & \text{if } \frac{1}{2}\left(1 - \frac{1}{n}\right) < x < \frac{1}{2}, \\ 1, & \text{if } \frac{1}{2} \le x \le 1. \end{cases}$$

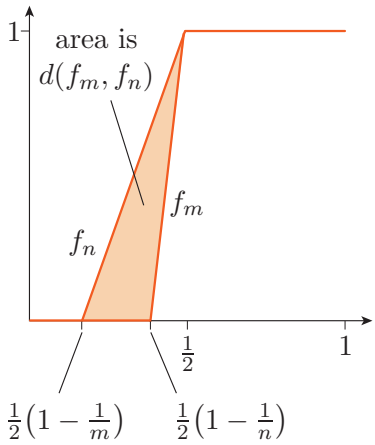

Figure 2.2

Either by integrating, or by finding the areas of appropriate triangles, we can easily show that if $N \in \mathbb{N}$ and $m > n > N$ (Figure 2.2), then

$$d(f_m, f_n) = \frac{1}{4}\left(\frac{1}{n} - \frac{1}{m}\right) < \frac{1}{2N}.$$

Hence (f_n) is a Cauchy sequence in $(C[0,1], d)$.

Now define $f\colon [0,1] \to \mathbb{R}$ by

$$f(x) = \begin{cases} 0, & \text{if } 0 \le x < \frac{1}{2}, \\ 1, & \text{if } \frac{1}{2} \le x \le 1. \end{cases}$$

Then,

$$\begin{aligned} \int_0^1 |f(x) - f_n(x)|\,dx &= \int_{\frac{1}{2}(1-\frac{1}{n})}^{\frac{1}{2}} 2n\left(x - \frac{1}{2}\left(1 - \frac{1}{n}\right)\right) dx \\ &= \left[n\left(x - \frac{1}{2}\left(1 - \frac{1}{n}\right)\right)^2\right]_{\frac{1}{2}(1-\frac{1}{n})}^{\frac{1}{2}} \\ &= n\left(\frac{1}{2n}\right)^2 = \frac{1}{4n}. \end{aligned}$$

Thus f is a candidate for the limit of the sequence (f_n). However, f is *not* an element of $C[0,1]$, since it is not continuous at the point $\frac{1}{2}$. In fact, for each $g \in C[0,1]$

$$\int_0^1 |g(x) - f(x)|\,dx > 0.$$

To see this, we first note that since g is continuous and f is not, we know that $f \neq g$. If f were continuous on $[0,1]$, we could then deduce that

$$d(f,g) = \int_0^1 |g(x) - f(x)|\,dx > 0.$$

Unfortunately, we cannot apply the metric d since f is not continuous and so not in $C[0,1]$. However, both f restricted to $[0, \frac{1}{2})$ and f restricted to $[\frac{1}{2}, 1]$ are continuous, and so we consider the metrics on $C[0, \frac{1}{2})$ and $C[\frac{1}{2}, 1]$ given respectively by

$$d_-(g,h) = \int_0^{\frac{1}{2}} |h(x) - g(x)|\,dx \text{ and } d_+(g,h) = \int_{\frac{1}{2}}^1 |h(x) - g(x)|\,dx.$$

The proof that d_- and d_+ are metrics is similar to the proof that d is a metric.

We know that there exists $x_0 \in [0,1]$ such that $f(x_0) \neq g(x_0)$ and so at least one of $d_-(f|_{[0,\frac{1}{2})}, g|_{[0,\frac{1}{2})})$ and $d_+(f|_{[\frac{1}{2},1]}, g|_{[\frac{1}{2},1]})$ is positive. Since

$$\int_0^1 |g(x) - f(x)|\,dx = d_-(f|_{[0,\frac{1}{2})}, g|_{[0,\frac{1}{2})}) + d_+(f|_{[\frac{1}{2},1]}, g|_{[\frac{1}{2},1]}),$$

we have $\int_0^1 |g(x) - f(x)| > 0$, as claimed.

We now use the Reverse Triangle Inequality to estimate

$$\begin{aligned} d(f_n, g) &= \int_0^1 |g(x) - f_n(x)|\, dx \\ &\geq \left| \int_0^1 |g(x) - f(x)|\, dx - \int_0^1 |f(x) - f_n(x)|\, dx \right| \\ &\geq \int_0^1 |g(x) - f(x)|\, dx - \int_0^1 |f(x) - f_n(x)|\, dx \\ &= \int_0^1 |g(x) - f(x)|\, dx - \frac{1}{4n}. \end{aligned}$$

But if we choose n large enough so that $\frac{1}{4n} < \frac{1}{2}\int_0^1 |g(x) - f(x)|\, dx$, then

$$d(f_n, g) > \tfrac{1}{2} \int_0^1 |g(x) - f(x)|\, dx > 0,$$

and so (f_n) does not converge to g.

Since g was an arbitrary function in $C[0,1]$, we conclude that the sequence (f_n) is not convergent and $(C[0,1], d)$ is not a complete metric space.

This is a little unfortunate in that there are many occasions when it is more natural to use the integral metric d rather than the $d_{\max}$ metric. This example shows that, in this case, it is not going to be enough to show that a sequence is Cauchy in order to deduce it is convergent.

This raises the question of whether we can find a complete metric space (X, d_X) that contains $(C[0,1], d)$ as a subspace. In fact, we can. In Section 4, we discuss a method for constructing the *smallest* complete space containing a given metric space as a subspace.

2.5 Completeness is not a homeomorphism invariant

Recall that in Chapter 21, we introduced the notion of homeomorphisms for metric spaces and saw that the property of being connected is a homeomorphism invariant: if (X, d) is connected and homeomorphic to (Y, e), then (Y, e) is connected. The following example shows that completeness is not a homeomorphism invariant.

Example 2.10

Let $\mathbb{R}$ have the Euclidean metric and let $X = (-1, 1)$ have the metric induced from the Euclidean metric for $\mathbb{R}$. The function $f\colon X \to \mathbb{R}$ given by $f(x) = \tan\left(\pi\frac{x}{2}\right)$ is a homeomorphism between $(-1, 1)$ and $\mathbb{R}$ (Figure 2.3 overleaf), and so $(X, d_X^{(1)})$ is homeomorphic to $(\mathbb{R}, d^{(1)})$.

We omit the proof that f is $(d_X^{(1)}, d^{(1)})$-continuous and f^{-1} is $(d^{(1)}, d_X^{(1)})$-continuous.

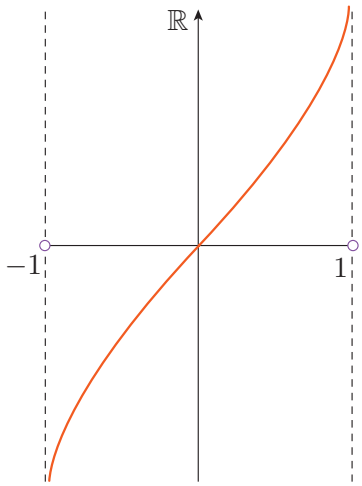

Figure 2.3

Now consider the sequence $\left(1-\frac{1}{n}\right)$ in X. It is a Cauchy sequence since, if $N \in \mathbb{N}$ and $n > m > N$, then

$$\left(1-\frac{1}{n}\right)-\left(1-\frac{1}{m}\right)=\frac{1}{m}-\frac{1}{n}<\frac{1}{N}\to\ 0 \text{ as } m,n\to\infty.$$

This sequence converges to 1 in $\mathbb{R}$, but does not converge in X, since $1 \notin (-1,1)$. Hence, $(X, d_X^{(1)})$ is incomplete.

But $(\mathbb{R}, d^{(1)})$ is complete and so we have two homeomorphic metric spaces, one complete and the other not: thus completeness is not a topological invariant.

We gain some insight into this example by noting that if we apply the homeomorphism f to the terms of the sequence $(1-\frac{1}{n})$ in $(-1,1)$, we obtain the image sequence $(\tan(\frac{\pi}{2}-\frac{\pi}{2n}))$ in $\mathbb{R}$. This is not a Cauchy sequence in $\mathbb{R}$; if it were, it would converge (since $\mathbb{R}$ is complete) and its limit would be $\tan\frac{\pi}{2}$, which is undefined.

Although completeness is not a homeomorphism invariant, it is a metric invariant:

if two metric spaces are metrically the same, then they are either both complete or both incomplete.

In order to make sense of this statement, we need to understand what it means for two metric spaces to be 'metrically the same'. The key property is the notion of distance: two metric spaces are metrically the same if we can map the points of one space to the other space, with a one–one and onto map, so that all distances are preserved. The following definition makes this precise.

Definition 2.11 *Isometric metric spaces*

Let (X, d_X) and (Y, d_Y) be metric spaces.

A function $f\colon X \to Y$ is an **isometry** if

1. f is one–one and onto
2. for each $x_1, x_2 \in X$, $d_X(x_1, x_2) = d_Y(f(x_1), f(x_2))$.

The metric spaces (X, d_X) and (Y, d_Y) are **isometric** if there is an isometry between X and Y.

Remarks

1. If $f\colon X \to Y$ is an isometry, then it is invertible (since it is one–one and onto) and its inverse, f^{-1}, is also an isometry.
2. Since an isometry preserves distances and is a bijection, a sequence in X converges if, and only if, its image under the isometry also converges in Y. Equivalently, a sequence in X is a d_X-Cauchy sequence if, and only if, its image under the isometry is a d_Y-Cauchy sequence in Y.

See Exercise 2.9.

3. Since closed sets and sequentially compact sets can be characterised in terms of sequences, an isometry preserves closed and (sequentially) compact sets, and thus also preserves open sets. This implies that connectedness is also preserved under isometry.

Exercise 2.8

Let d be the Euclidean metric for $\mathbb{R}$. Show that:

(a) $([0,1], d_{[0,1]})$ is isometric to $([1,2], d_{[1,2]})$

(b) $([0,1], d_{[0,1]})$ is not isometric to $([0,2], d_{[0,2]})$.

Exercise 2.9

Let (X, d_X) and (Y, d_Y) be metric spaces, and let $f\colon X \to Y$ be an isometry. Show that if (a_n) is a sequence in X that is d_X-convergent to a, then $(f(a_n))$ is d_Y-convergent to $f(a)$ in Y.

Exercise 2.8(b) gives an example that shows that homeomorphic metric spaces need not be isometric. Thus the following theorem does not contradict Example 2.10.

Theorem 2.12

Let (X, d_X) and (Y, d_Y) be isometric metric spaces. If (X, d_X) is complete, then (Y, d_Y) is also complete.

Proof Suppose that (X, d_X) and (Y, d_Y) are isometric metric spaces and let $f\colon X \to Y$ be an isometry. Let (X, d_X) be complete, and let (a_n) be a d_Y-Cauchy sequence in Y. Then $(f^{-1}(a_n))$ is a d_X-Cauchy sequence in X since, for any $n, m \in \mathbb{N}$,

Since f is an isometry, it is invertible.

$$d_X(f^{-1}(a_n), f^{-1}(a_m)) = d_Y(a_n, a_m).$$

Since (X, d_X) is complete, there is an $x \in X$ such that $(f^{-1}(a_n))$ is d_X-convergent to x. But

$$d_Y(a_n, f(x)) = d_X(f^{-1}(a_n), x) \to 0, \text{ as } n \to \infty,$$

so (a_n) is d_Y-convergent to $f(x)$ in Y. Hence, (Y, d_Y) is complete. ■

Isometries play a big role in Section 4, where we discuss how to form the *completion* of a metric space – the smallest complete metric space that contains a given space.

3 The Contraction Mapping Theorem

In this section, we show how complete metric spaces play an important role in the theory of fixed points of mappings. Finding the fixed points of a particular function – that is, those points that are mapped to themselves – is frequently an important part of solving a mathematical problem.

We give a formal definition of fixed points in Definition 3.1.

3.1 Fixed points and the Contraction Mapping Theorem

Many problems in mathematics can be written in the form:

solve the equation $f(x) = x$.

Such points X are known as *fixed points* of f.

Definition 3.1 *Fixed point*

Let X be a set and let $T\colon X \to X$ be a function.

A **fixed point** of T is a point $x \in X$ such that $T(x) = x$.

We use T to denote the function and not f because the set X may have functions as its elements.

We shall see some examples shortly.

Fixed points are surprisingly important and there are many theorems that provide conditions that guarantee a function T has a fixed point. In this section, we are going to prove one such result, the Contraction Mapping Theorem. It ensures that, under appropriate conditions, a function has a *unique* fixed point. This theorem also provides a method for finding an approximate value for the fixed point.

Exercise 3.1

Let $f\colon \mathbb{R} \to \mathbb{R}$ be a function and define $g\colon \mathbb{R} \to \mathbb{R}$ by $g(x) = x - f(x)$. Show that X is a fixed point of g if, and only if, X is a zero of f – that is, $f(X) = 0$.

Exercise 3.2

Let $f\colon \mathbb{R} \to \mathbb{R}$ be a function and for $r \in \mathbb{R} - \{0\}$ define $g_r\colon \mathbb{R} \to \mathbb{R}$ by $g_r(x) = x - rf(x)$. Show that X is a fixed point of g_r if, and only if, X is a zero of f.

Thus the problem of finding the zeros of a real-valued function can be converted into a problem concerning the fixed points of a (different) real-valued function. For example, finding the zeros of

$\cos x - \exp x$, for $x \in \mathbb{R}$ (see Figure 3.1),

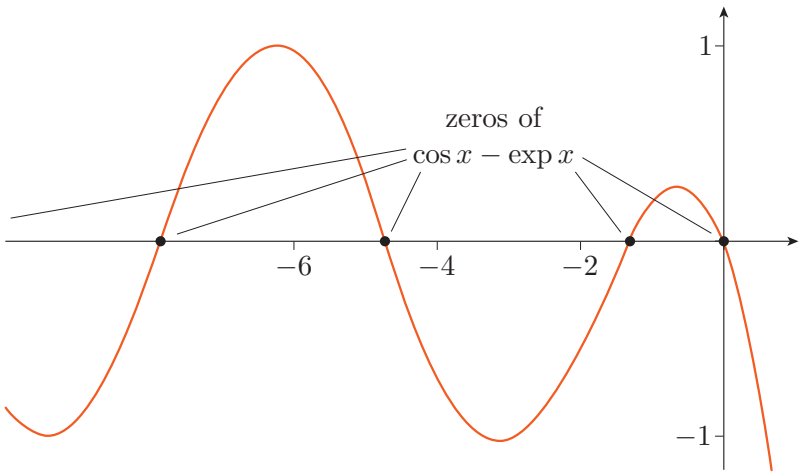

Figure 3.1

is the same as finding the fixed points of the function $g\colon \mathbb{R} \to \mathbb{R}$ given by

$$g(x) = x - \cos x + \exp x \quad \text{(see Figure 3.2).}$$

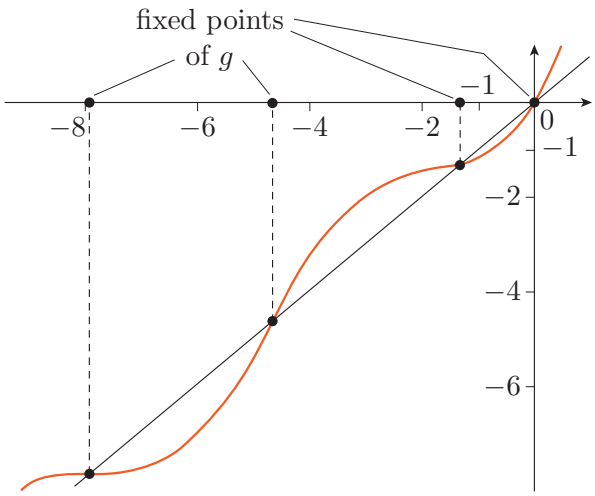

Figure 3.2

Our next example of a fixed-point problem involves a mapping defined on $C[0,1]$.

Example 3.2

Suppose that a function $f \in C[0,1]$ satisfies the differential equation

$$\frac{df}{ds} = sf(s), \quad f(0) = 1 \quad (0 \le s \le 1). \tag{1}$$

Integrating both sides of this equation with respect to s gives

$$\int_0^t \frac{df}{ds}\, ds = \int_0^t sf(s)\, ds$$

and so

$$f(t) - f(0) = \int_0^t sf(s)\,ds, \text{ for } 0 \le t \le 1.$$

Using the boundary condition $f(0) = 1$, we obtain an *integral equation* that any solution f of the differential equation must satisfy; namely,

$$f(t) = 1 + \int_0^t sf(s)\,ds, \text{ for } 0 \le t \le 1.$$

Let us define a map T on $C[0,1]$, by

$$T(f)(t) = 1 + \int_0^t sf(s)\,ds, \text{ for } f \in C[0,1]. \tag{2}$$

Then $T(f)$ is itself a function from $[0,1]$ to $\mathbb{R}$ and any solution f of the differential equation (1) is also a fixed point of the map T, $T(f) = f$.

We now show that T is a map from $C[0,1]$ to itself. First note that $s \mapsto sf(s)$ is continuous, since it is the product of continuous functions, and so belongs to $C[0,1]$. Also, the map F defined by

$$F(g)(t) = \int_0^t g(s)ds, \text{ for } 0 \le t \le 1,$$

See Book D, Chapter 15, Worked Exercise 5.7.

is a map from $C[0,1]$ to itself. So the function G defined by

$$G(f)(t) = \int_0^t sf(s)ds, \text{ for } 0 \le t \le 1,$$

is in $C[0,1]$. The constant function 1 is also in $C[0,1]$ and so T is a map from $C[0,1]$ to itself.

This example shows us that any solution of the differential equation (1) is a fixed point of the mapping T given in (2). Thus finding all the fixed points of the mapping T also gives us all the potential solutions to the differential equation. (It may be the case that the integral equation has fixed points that are not solutions of the differential equation – we will give criteria later that ensure this does not occur.)

Exercise 3.3

Consider the differential equation

$$\frac{df}{ds} = \tfrac{1}{2}\cos f(s), \quad f(0) = 0 \quad (0 \le s \le 1).$$

By integrating both sides of this equation, find an integral equation that every solution of this differential equation must satisfy.

Hence define a map $T\colon C[0,1] \to C[0,1]$ such that any solution of the differential equation is also a fixed point of this map.

The theorem we now prove is the Contraction Mapping Theorem. This theorem guarantees the existence of a fixed point when the mapping from a complete metric space to itself is a special type of map known as a *contraction mapping*.

Definition 3.3 *Contraction mapping*

Let (X, d) be a metric space. A function $T\colon X \to X$ is a **contraction mapping** for d if there is a real number $0 \leq \lambda < 1$ such that

$$d(T(x), T(y)) \leq \lambda d(x, y) \text{ for each } x, y \in X \quad \text{(Figure 3.3)}.$$

Any such real number λ is known as a **contraction ratio** for T.

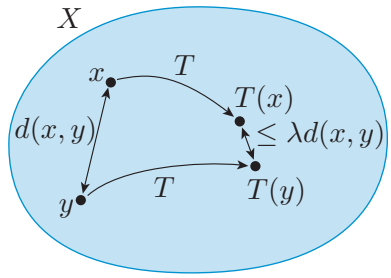

Figure 3.3

Remarks

1. We usually suppress mention of the metric and just say a map is a contraction mapping.
2. A contraction mapping is a Lipschitz function from a set X to itself with Lipschitz constant *strictly less than* 1. For example, $T\colon \mathbb{R} \to \mathbb{R}$ given by $T(x) = \frac{1}{2}x$ is a Lipschitz mapping with Lipschitz constant $\frac{1}{2}$ for the Euclidean metric on $\mathbb{R}$, and so is a contraction mapping.
3. A contraction mapping has many different possible contraction ratios: for if $\lambda \in (0, 1)$ is a contraction ratio, then so is any μ for which $\lambda < \mu < 1$.

Worked Exercise 3.4

Show that for the metric $d_{\max}$ on $C[0, 1]$, the map $T\colon C[0, 1] \to C[0, 1]$ defined by

$$T(f)(t) = 1 + \int_0^t s f(s) ds, \text{ for } 0 \leq t \leq 1$$

is a contraction mapping with contraction ratio $\frac{1}{2}$.

You met this map in Example 3.2.

Solution

We have already seen in Example 3.2 that T is a map from $C[0, 1]$ to $C[0, 1]$, and so it only remains to show that it is a contraction.

Given $f, g \in C[0, 1]$ and $t \in [0, 1]$, we have

$$\begin{aligned} |T(f)(t) - T(g)(t)| &= \left| \int_0^t s(f(s) - g(s))\, ds \right| \\ &\leq \int_0^t |s(f(s) - g(s))|\, ds \\ &\leq d_{\max}(f, g) \int_0^t s\, ds \\ &\leq d_{\max}(f, g) \int_0^1 s\, ds = \tfrac{1}{2} d_{\max}(f, g). \end{aligned}$$

Since t is arbitrary, we deduce that

$$d_{\max}(T(f), T(g)) \leq \tfrac{1}{2} d_{\max}(f, g),$$

and so T is a contraction mapping with contraction ratio $\frac{1}{2}$.

Exercise 3.4

Let (X, d) be a metric space and let $T\colon X \to X$ be a contraction mapping for d with contraction ratio 0. Show that T is a *constant* function.

Exercise 3.5

You met this map in Exercise 3.3.

Let $T\colon C[0,1] \to C[0,1]$ be defined by

$$T(f)(t) = \int_0^t \tfrac{1}{2} \cos f(s)\, ds, \text{ for } 0 \leq t \leq 1.$$

Show that T is a contraction mapping with contraction ratio $\frac{1}{2}$.

Hint: use the fact that $|\cos a - \cos b| \leq |a - b|$ for any real numbers a and b.

The next lemma shows that contraction mappings can have at most one fixed point. We will use this result later in the proof of the Contraction Mapping Theorem.

Lemma 3.5

Let (X, d) be a metric space. Suppose that $T\colon X \to X$ is a contraction mapping for d. Then T has at most one fixed point.

Proof Suppose that $x, y \in X$ are both fixed points of T – that is,

$$T(x) = x \text{ and } T(y) = y.$$

Since T is a contraction mapping,

$$d(x, y) = d(T(x), T(y)) \leq \lambda d(x, y),$$

where $0 \leq \lambda < 1$. Rearranging, we obtain

$$(1 - \lambda) d(x, y) \leq 0,$$

which implies that $d(x, y) \leq 0$, since $0 \leq \lambda < 1$. But $d(x, y) \geq 0$, and so $d(x, y) = 0$. Hence, since d is a metric, $x = y$. Thus T has at most one fixed point. ■

Remark

The conclusion of this lemma can fail if $T\colon X \to X$ is a Lipschitz map with Lipschitz constant 1 – that is if,

$$d(T(x), T(y)) \leq d(x, y) \text{ for } x, y \in X.$$

For example, consider $T\colon \mathbb{R} \to \mathbb{R}$ (with the Euclidean metric) where $T(x) = x$. For this mapping, *every* real number is a fixed point. This is why we require a contraction mapping to have a Lipschitz constant that is strictly less than one – it guarantees that it can have at most one fixed point.

Exercise 3.6

Let $(0, 1]$ have metric induced from the Euclidean metric for $\mathbb{R}$, and define $T\colon (0, 1] \to (0, 1]$ by $T(x) = \frac{1}{2}x$. Show that:

(a) T is a contraction mapping with contraction ratio $\frac{1}{2}$

(b) T has no fixed points.

Exercise 3.6 shows that a contraction mapping need not have any fixed points: the difficulty is caused by the fact that the point that 'should be' the fixed point of the mapping is not in the space. This difficulty will not arise for a *complete* metric space.

Theorem 3.6 *Contraction Mapping Theorem*

Let (X, d) be a complete metric space, and let $T\colon X \to X$ be a contraction mapping for d. Then there is a unique element $x_T \in X$ such that

$$T(x_T) = x_T.$$

Moreover, for each $x \in X$, the sequence of iterates

$$x,\ T(x),\ T(T(x)),\ T(T(T(x))),\ \ldots$$

converges to this unique fixed point x_T; that is, if

$$T^n(x) = (T \circ T \circ \cdots \circ T)(x) \ (T \text{ composed with itself } n \text{ times}),$$

then for each $x \in X$

$$d(T^n(x), x_T) \to 0 \quad \text{as } n \to \infty.$$

This is also known as the Banach Fixed Point Theorem after the Polish mathematician Stefan Banach (1892–1945). Banach was a prolific analyst and has many concepts and theorems named after him.

Proof Our proof involves three steps:

1. We show that for each $x \in X$, the sequence $x, T(x), T^2(x), \ldots, T^n(x), \ldots$ is a Cauchy sequence for d. Since (X, d) is complete, this sequence converges to some element $x_T \in X$.
2. We verify that $T(x_T) = x_T$ – that is, x_T is a fixed point of T.
3. We deduce from Lemma 3.5 that the fixed point is unique.

Since T is a contraction mapping for d, there is a number λ (with $0 \leq \lambda < 1$) such that

$$d(T(x), T(y)) \leq \lambda d(x, y) \text{ for each } x, y \in X.$$

Now fix $x \in X$ and consider the sequence of iterates: $x, T(x), T^2(x), \ldots$.

Show that $(T^n(x))$ is a Cauchy sequence

We need to show that for each $\varepsilon > 0$, there is an $N \in \mathbb{N}$ such that

$$d(T^n(x), T^m(x)) < \varepsilon \text{ for each } n, m > N.$$

So let $\varepsilon > 0$ be given.

We start by estimating $d(T^k(x), T^{k+1}(x))$, the distance between two successive terms of the sequence, for $k = 1, 2, \ldots$:

$$\begin{aligned} d(T^k(x), T^{k+1}(x)) &= d(T(T^{k-1}(x)), T(T^k(x))) \\ &\leq \lambda d(T^{k-1}(x), T^k(x)), \text{ since } T \text{ is a contraction mapping} \\ &= \lambda d(T(T^{k-2}(x)), T(T^{k-1}(x))) \\ &\leq \lambda^2 d(T^{k-2}(x), T^{k-1}(x)), \text{ since } T \text{ is a contraction mapping} \\ &\vdots \quad \text{(repeating this } k \text{ times)} \\ &\leq \lambda^k d(x, T(x)). \end{aligned} \tag{3}$$

We define $T^0(x) = x$.

Let $n, m \in \mathbb{N}$ with $m > n$. Then, by (3) and repeated use of the Triangle Inequality for d, we find:

$$\begin{aligned} &d(T^n(x), T^m(x)) \\ &\leq d(T^n(x), T^{n+1}(x)) + d(T^{n+1}(x), T^{n+2}(x)) + \cdots + d(T^{m-1}(x), T^m(x)) \\ &\leq \lambda^n d(x, T(x)) + \lambda^{n+1} d(x, T(x)) + \lambda^{n+2} d(x, T(x)) + \cdots + \lambda^{m-1} d(x, T(x)) \\ &= (\lambda^n + \lambda^{n+1} + \lambda^{n+2} + \cdots + \lambda^{m-1}) d(x, T(x)) \\ &= \frac{\lambda^n - \lambda^m}{1 - \lambda} d(x, T(x)) \\ &\leq \frac{\lambda^n}{1 - \lambda} d(x, T(x)). \end{aligned} \tag{4}$$

Since $0 \leq \lambda < 1$ and x is fixed, we can choose $N \in \mathbb{N}$ so that

$$\frac{\lambda^N}{1 - \lambda} d(x, T(x)) < \varepsilon.$$

Hence for each $m > n > N$,

$$d(T^n(x), T^m(x)) \leq \frac{\lambda^n}{1 - \lambda} d(x, T(x)) \leq \frac{\lambda^N}{1 - \lambda} d(x, T(x)) < \varepsilon,$$

since $\lambda^n \leq \lambda^N$ for $0 \leq \lambda < 1$. Thus the sequence is a Cauchy sequence. But (X, d) is a *complete* metric space and so the sequence $(T^n(x))$ is d-convergent to some point $x_T \in X$.

Verify that x_T is a fixed point

By the Triangle Inequality for d, we have for each $n \geq 1$,

$$\begin{aligned} 0 \leq d(x_T, T(x_T)) &\leq d(x_T, T^n(x)) + d(T^n(x), T(x_T)) \\ &\leq d(x_T, T^n(x)) + \lambda d(x_T, T^{n-1}(x)). \end{aligned}$$

Since n is arbitrary and both $d(x_T, T^{n-1}(x))$ and $d(x_T, T^n(x)) \to 0$ as $n \to \infty$, we can take the limit as $n \to \infty$ on the right-hand side. We

conclude that

$0 \leq d(x_T, T(x_T)) \leq 0 + \lambda \cdot 0 = 0$

– that is, $d(x_T, T(x_T)) = 0$ and so $T(x_T) = x_T$.

Show that the fixed point is unique

This follows immediately from Lemma 3.5, which tells us that there is at most one fixed point. ■

Worked Exercise 3.7

Use the Contraction Mapping Theorem to show that the equation $x = \frac{1}{2}\cos x$ has a unique real solution.

We could also use the Intermediate Value Theorem to show that a root exists, but that would not show that there is exactly one root.

This is common in situations involving the Contraction Mapping Theorem.

Solution

This exercise does not specify which metric space we should work in. We choose a complete metric space and a contraction mapping T appropriate to the problem. In this case, there is a natural choice of mapping and space.

Let $T\colon \mathbb{R} \to \mathbb{R}$ be given by $T(x) = \frac{1}{2}\cos x$, and consider the complete metric space $(\mathbb{R}, d^{(1)})$.

Note that $T(x) = x$ if, and only if, $x = \frac{1}{2}\cos x$, and so the roots of this equation correspond exactly to the fixed points of the mapping T.

It remains only to check that T is a contraction. For $x, y \in \mathbb{R}$ with $x \neq y$, the Mean Value Theorem implies that there is $c \in (x, y)$ such that

$$\left|\frac{T(x) - T(y)}{x - y}\right| = \left|\frac{\frac{1}{2}(\cos x - \cos y)}{x - y}\right| = |T'(c)|.$$

See the *Handbook*.

Now

$$|T'(c)| = \left|\tfrac{1}{2}\sin c\right| \leq \tfrac{1}{2}, \text{ for } c \in \mathbb{R},$$

and so T is a contraction for the Euclidean metric with contraction ratio $\frac{1}{2}$. Hence we deduce from the Contraction Mapping Theorem that T has a unique fixed point, and so the original equation also has a unique real solution.

Exercise 3.7

Let $[1, \infty)$ have the induced Euclidean metric and define $T\colon [1, \infty) \to [1, \infty)$ by

$$T(x) = x + \tfrac{1}{x}.$$

(a) Show that

$$|T(x) - T(y)| < |x - y| \text{ for each } x, y \in [1, \infty).$$

(b) Does this map have a fixed point?

(c) Does this contradict the Contraction Mapping Theorem?

By looking at the proof of the Contraction Mapping Theorem, we can obtain an estimate of how far $T^n(x)$ is from x_T (the unique fixed point) in terms of $d(x, T(x))$. This is useful when using a computer to find an approximate value for the fixed point.

Corollary 3.8

Let (X, d) be a complete metric space and let $T\colon X \to X$ be a contraction mapping for d with contraction ratio $\lambda \in [0, 1)$. Then for each $x \in X$ and each $n \in \mathbb{N}$,

$$d(x_T, T^n(x)) \leq \frac{\lambda^n}{1-\lambda}\, d(x, T(x)),$$

where x_T is the fixed point of the map T.

Proof Let $x \in X$ and take $m, n \in \mathbb{N}$ with $m > n$. By the Triangle Inequality for d,

$$\begin{aligned} d(x_T, T^n(x)) &\leq d(x_T, T^m(x)) + d(T^m(x), T^n(x)) \\ &\leq d(x_T, T^m(x)) + \frac{\lambda^n}{1-\lambda}\, d(x, T(x)), \quad \text{by (4)}. \end{aligned}$$

But by the Contraction Mapping Theorem, $d(x_T, T^m(x)) \to 0$ as $m \to \infty$, and so, since m can be arbitrarily large, we deduce that

$$d(x_T, T^n(x)) \leq \frac{\lambda^n}{1-\lambda}\, d(x, T(x)).$$

■

Consider the mapping $T\colon \mathbb{R} \to \mathbb{R}$ given by $T(x) = \frac{1}{2}\cos x$. In Worked Exercise 3.7 we showed that T is a contraction with contraction ratio $\frac{1}{2}$ and so, by the Contraction Mapping Theorem, T has a unique fixed point x_T. We can now use Corollary 3.8 to estimate the value of x_T.

Taking $n = 1$ and $x = 0$, we have

$$|x_T - T(0)| \leq \frac{\frac{1}{2}}{1-\frac{1}{2}}|0 - T(0)| = \tfrac{1}{2}.$$

Since $T(0) = \frac{1}{2}$, the fixed point of T (which is the root of $x = \frac{1}{2}\cos x$) lies somewhere in $[0, 1]$. We can use Corollary 3.8 to obtain finer approximations to x_T by looking at $T^n(0)$ for larger values of n. For example, taking $n = 2$, we obtain

$$|x_T - T^2(0)| \leq \frac{(\frac{1}{2})^2}{1-\frac{1}{2}}|0 - T(0)| = \tfrac{1}{4}.$$

Since $T^2(0) = \frac{1}{2}\cos\frac{1}{2}$, this implies that x_T lies in the interval

$$[\tfrac{1}{2}\cos\tfrac{1}{2} - \tfrac{1}{4}, \tfrac{1}{2}\cos\tfrac{1}{2} + \tfrac{1}{4}] \subseteq [0.188, 0.689].$$

By increasing n, we can find smaller and smaller intervals that contain the fixed point.

Exercise 3.8

Consider the mapping $T\colon \mathbb{R} \to \mathbb{R}$ given by $T(x) = \frac{1}{2}\cos x$. Find an interval of length at most 0.1 containing the unique fixed point of T.

3.2 Applications of the Contraction Mapping Theorem

We now give some applications of the Contraction Mapping Theorem.

In the next chapter, we use the Contraction Mapping Theorem to define fractals.

Locating the zeros of real functions

If a continuous real-valued function on an interval is a contraction mapping, then we can use the Contraction Mapping Theorem to deduce the existence of a unique fixed point in that interval, and Corollary 3.8 to estimate how far we are from this fixed point when we iterate an initial guess.

Unfortunately, many of the functions whose zeros we wish to find cannot be directly rearranged to give a fixed point problem involving contractions. For example, if we wish to find the zeros of

$$x^2 - e^{-x}$$

(Figure 3.4) then, following Exercise 3.1, it is natural to look for the fixed points of the function $g\colon \mathbb{R} \to \mathbb{R}$ given by $g(x) = x - x^2 + e^{-x}$. But g is not a contraction mapping, and so we cannot apply the Contraction Mapping Theorem *directly*. However, by carefully defining an auxiliary function g, we *can* use the Contraction Mapping Theorem to find good approximations to the zeros of f.

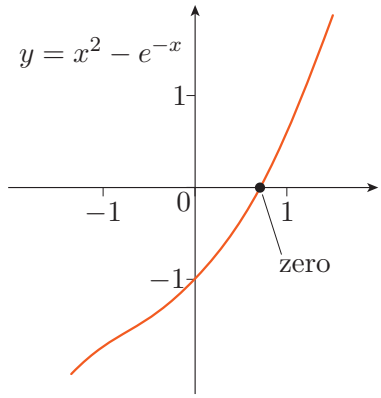

Figure 3.4

A sketch of the graph of f suggests that there is only one zero.

So suppose that we wish to find a zero of $f\colon \mathbb{R} \to \mathbb{R}$, given by $f(x) = x^2 - e^{-x}$. The function f is continuous, $f(0) = -1$ and $f(1) = 1 - e^{-1} > 0$. By the Intermediate Value Theorem, there is a zero of f in the interval $[0, 1]$.

The idea is to notice that the zeros of f correspond to the fixed points of the function g_r given by $g_r(x) = x - rf(x)$ (see Exercise 3.2) for *any* non-zero real number r. Thus, if we can find a non-zero real number r and a closed interval I such that $g_r\colon I \to I$ and g_r is a contraction mapping on this interval, then we can deduce from the Contraction Mapping Theorem (using the induced Euclidean metric for I) that f has a unique zero in this interval.

With this objective in mind, we estimate the derivative of g_r on the interval $[0, 1]$, where f has a zero. For f, we have $f'(x) = 2x + e^{-x}$, and so

$$1 \leq f'(x) \leq 3, \text{ for each } x \in [0, 1] \quad \text{(Figure 3.5 overleaf)}.$$

This tells us that f is strictly increasing on the interval $[0, 1]$.

This method tells us that there is a unique zero in the interval I but it does not tell us whether there are any zeros outside this interval. We need to use other methods to decide this.

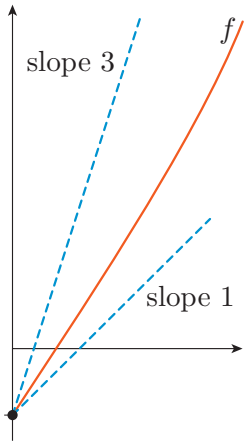

Figure 3.5

The function g_r has derivative

$$g_r'(x) = 1 - rf'(x), \text{ for all } x \in [0,1].$$

Hence for each $r > 0$,

$$1 - 3r \leq g_r'(x) \leq 1 - r \text{ whenever } 0 \leq x \leq 1.$$

Moreover, if $r \leq \frac{1}{3}$, then $1 - 3r \geq 0$ and so g_r is non-decreasing on $[0,1]$. Hence, for $0 \leq x \leq 1$ and $0 < r \leq \frac{1}{3}$,

$$g_r(x) \geq g_r(0) = 0 - rf(0) = r > 0$$

and

$$g_r(x) \leq g(1) = 1 - rf(1) = 1 - r\left(1 - \tfrac{1}{e}\right) < 1,$$

and so $g_r \colon [0,1] \to [0,1]$. Moreover, for $0 < r \leq \frac{1}{3}$, it follows from the Mean Value Theorem that for each $x, y \in [0,1]$,

$$|g_r(x) - g_r(y)| \leq \max_{0 \leq c \leq 1} |g_r'(c)||x - y| \leq (1 - r)|x - y|.$$

Thus, for $0 < r \leq \frac{1}{3}$, g_r is a contraction mapping on $[0,1]$ with contraction ratio $1 - r$. By the Contraction Mapping Theorem, g_r has a unique fixed point in the interval $[0,1]$ and thus f has a unique zero in $[0,1]$. If $r = \frac{1}{3}$ and x_r is the fixed point of g_r (and the zero of f) then, by Corollary 3.8,

$$|x_r - g_r^n(x_0)| \leq \frac{(1-r)^n}{1-(1-r)}|x_0 - g_r(x_0)| = 3\left(\tfrac{2}{3}\right)^n |x_0 - g_r(x_0)| \leq 3\left(\tfrac{2}{3}\right)^n,$$

whenever $x_0 \in [0,1]$ is an initial guess for the value of the fixed point. This method can be used to find a zero of any differentiable function f that is strictly monotonic on an open interval containing the zero being sought.

Theorem 3.9

Let $a < b \in \mathbb{R}$, and suppose that $f \colon [a,b] \to \mathbb{R}$ is continuous and differentiable on $[a,b]$. If

1. $f(a) < 0 < f(b)$
2. there are $m, M \in \mathbb{R}$ with $0 < m \leq M$ such that

$$m \leq f'(x) \leq M \text{ for each } x \in [a,b],$$

then for $0 < r \leq \frac{1}{M}$, the function $g_r \colon [a,b] \to \mathbb{R}$ given by $g_r(x) = x - rf(x)$ maps $[a,b]$ into itself and is a contraction mapping of $[a,b]$ with contraction ratio $1 - mr$.

It then follows from the Contraction Mapping Theorem (using the induced Euclidean metric for $[a,b]$) that a function f satisfying the conditions of Theorem 3.9 has a unique zero in $[a,b]$. We can then use Corollary 3.8 (applied to g_r) to find an approximate location for this zero to any desired accuracy.

Exercise 3.9

Prove Theorem 3.9.

Solving differential equations

Recall from Example 3.2 that any solution of the differential equation

$$\frac{df}{ds} = sf(s), \quad f(0) = 1 \quad (0 \leq s \leq 1)$$

must also satisfy the integral equation

$$f(t) = 1 + \int_0^t sf(s)\,ds \text{ for } 0 \leq t \leq 1.$$

We wish to find the fixed points of the map $T\colon C[0,1] \to C[0,1]$, given by

$$T(f)(t) = 1 + \int_0^t sf(s)\,ds.$$

In Worked Exercise 3.4, we saw that T is a contraction mapping for $d_{\max}$ with contraction ratio $\frac{1}{2}$. In Subsection 2.4, we proved that $(C[0,1], d_{\max})$ is a complete metric space. So, by the Contraction Mapping Theorem, there is a unique function $f \in C[0,1]$ for which $f = T(f)$ – that is,

$$f(t) = 1 + \int_0^t sf(s)\,ds. \tag{5}$$

And we can find good approximations to the fixed point of T by using Corollary 3.8.

We start by calculating $T(f_0)$ for a simple function $f_0 \in C[0,1]$ – say $f_0(t) = 1$ for each $t \in [0,1]$, so that

$$T(f_0)(t) = 1 + \int_0^t s\,ds = 1 + \tfrac{1}{2}t^2.$$

Hence,

$$|T(f_0)(t) - f_0(t)| = \tfrac{1}{2}t^2,$$

and so

$$d_{\max}(T(f_0), f_0) = \max\left\{\tfrac{1}{2}t^2 : t \in [0,1]\right\} = \tfrac{1}{2}.$$

Thus if $f \in C[0,1]$ is the fixed point of T, then by Corollary 3.8,

$$d_{\max}(f, T^n(f_0)) \leq \frac{(\frac{1}{2})^n}{1 - \frac{1}{2}} d_{\max}(T(f_0), f_0) \leq (\tfrac{1}{2})^n.$$

We now calculate $T^2(f_0)$ for $t \in [0,1]$: we have

$$\begin{aligned} T^2(f_0)(t) &= T(T(f_0))(t) \\ &= 1 + \int_0^t s(1 + \tfrac{1}{2}s^2)\,ds \\ &= 1 + \tfrac{1}{2}t^2 + \tfrac{1}{8}t^4 = 1 + \left(\tfrac{1}{2}t^2\right) + \tfrac{1}{2}\left(\tfrac{1}{2}t^2\right)^2. \end{aligned}$$

Similarly, we can find $T^3(f_0)$: for $t \in [0,1]$, we have

$$\begin{aligned} T^3(f_0)(t) &= 1 + \int_0^t s\left(1 + \left(\tfrac{1}{2}s^2\right) + \tfrac{1}{2}\left(\tfrac{1}{2}s^2\right)^2\right) ds \\ &= 1 + \left(\tfrac{1}{2}t^2\right) + \tfrac{1}{2!}\left(\tfrac{1}{2}t^2\right)^2 + \tfrac{1}{3!}\left(\tfrac{1}{2}t^2\right)^3. \end{aligned}$$

In fact, on noting that

$$\int_0^t s\left(\tfrac{1}{2}s^2\right)^k ds = \frac{t^{2(k+1)}}{2^k(2k+2)} = \frac{1}{k+1}\left(\tfrac{1}{2}t^2\right)^{k+1},$$

we can use mathematical induction to conclude that

$$T^n(f_0)(t) = \sum_{k=0}^{n} \frac{1}{k!}\left(\tfrac{1}{2}t^2\right)^k.$$

Since convergence for the $d_{\max}$ metric implies pointwise convergence, the fixed point f of T is given by

$$f(t) = \lim_{n\to\infty}\left(T^n f_0\right)(t) = \sum_{k=0}^{\infty} \frac{1}{k!}\left(\tfrac{1}{2}t^2\right)^k = \exp(\tfrac{1}{2}t^2).$$

Thus $f(t) = \exp(\frac{1}{2}t^2)$ is the fixed point of the map T.

Exercise 3.10

Let $f(t) = \exp(\frac{1}{2}t^2)$. Verify that f satisfies (5).

We have found the (unique) fixed point of T in $(C[0,1], d_{\max})$, but it still remains to verify that this function also satisfies our original differential equation

$$\frac{df}{ds} = sf(s), \quad f(0) = 1 \quad (0 \le s \le 1).$$

The following result, whose proof we omit, guarantees that a solution to the integral equation corresponding to a given differential equation is also a solution of the original differential equation.

Theorem 3.10

Suppose that $F\colon \mathbb{R}^2 \to \mathbb{R}$ is a $(d^{(2)}, d^{(1)})$-continuous function and $a \in \mathbb{R}$ is given. Then $f \in C[0,1]$ is a solution of the differential equation

$$\frac{df}{ds} = F(s, f(s)), \quad f(0) = a \quad (0 \le s \le 1)$$

if, and only if,

$$f(t) = a + \int_0^t F(s, f(s))\, ds \quad \text{for each } 0 \le t \le 1.$$

Exercise 3.11

(a) Use the results of Exercises 3.3 and 3.5 to show that the differential equation

$$\frac{df}{ds} = \tfrac{1}{2}\cos f(s), \quad f(0) = 0 \quad (0 \leq s \leq 1)$$

has a unique solution in $C[0,1]$.

(b) By taking $f_0(s) = 0$ for $0 \leq s \leq 1$ as your initial approximation to this solution f, find an estimate $g \in C[0,1]$ of f for which

$$d_{\max}(f,g) \leq \tfrac{1}{4}.$$

Hint: you will need to use Theorem 3.10 to show that you have a unique solution to the differential equation.

4 Completion

In Section 3 we saw that, in complete metric spaces, contraction mappings have a unique fixed point, and that this is useful for solving a range of mathematical problems. In this section we show you how to find the 'smallest' complete metric space that contains a given (presumably incomplete) metric space.

4.1 Defining the completion

Suppose that we are given a metric space (X,d) that is incomplete, that is it has at least one divergent Cauchy sequence. Is it always possible to find a complete metric space that contains (X,d) as a metric subspace? This is an important question, because if it is not possible, then the utility of the Contraction Mapping Theorem is restricted.

In Section 1, we saw that $((0,1], d^{(1)}_{(0,1]})$ is not a complete metric space, but that $(\mathbb{R}, d^{(1)})$ is complete and contains $((0,1], d^{(1)}_{(0,1]})$. But $([0,1], d^{(1)}_{[0,1]})$ also contains $((0,1], d^{(1)})$ as a metric subspace and $([0,1], d^{(1)}_{[0,1]})$ is a complete metric space by Proposition 2.2, since $[0,1]$ is a closed subset of $\mathbb{R}$. Thus we have several possible choices for a complete metric space containing $((0,1], d^{(1)})$, with $([0,1], d^{(1)}_{[0,1]})$ appearing to be the 'smallest'.

Given a metric space containing (X,d), we aim to identify the 'smallest' complete metric space that contains (X,d) as a subspace. We do this by looking for a complete metric space (X^*, d^*) for which (X,d) is a *dense* subspace – that is, $\mathrm{Cl}_{(X^*,d^*)}(X) = X^*$ and $d = d^*_X$ is the induced metric from d^* on X. In particular, each point $x \in X^*$ is a d^*-closure point of X and so $B_{d^*}(x,r) \cap X \neq \emptyset$ for each $r > 0$.

Strictly, we should indicate that these are metrics induced from the Euclidean metric for $\mathbb{R}$, but this makes the notation cumbersome and so we don't bother.

In the case of $((0,1], d^{(1)})$, it is straightforward to identify such a metric space (X^*, d^*): it is $([0,1], d^{(1)})$.

Exercise 4.1

Find a complete metric space (X^*, d^*) for which $(\mathbb{Q}, d^{(1)})$ is a dense subspace.

If we can find a complete metric space (Y, e) that contains (X, d) as a subspace (so $X \subseteq Y$ and d is the induced metric, e_X, on X), then we take $X^* = \mathrm{Cl}_{(Y,e)}(X)$ and let d^* be the induced metric (from e) on X^* – Proposition 2.2 then tells us that (X^*, d^*) is also a complete metric space. Moreover, by definition it is a complete metric space that contains (X, d) as a dense set and it is the *smallest* such space: for if $y \in \mathrm{Cl}_{(Y,e)}(X) - X = X^* - X$ is omitted, then there is a sequence in X that e-converges to y in Y (and hence is e-Cauchy) but does not converge in $X^* - \{y\}$.

A sequence (x_n), all of whose terms lie in X, is d-Cauchy if, and only if, it is e-Cauchy when considered as a sequence in Y.

For each of the metric spaces $((0,1], d^{(1)})$ and $(\mathbb{Q}, d^{(1)})$, it is easy to find a complete metric space that contains it as a dense subspace, but in general it can be much harder to do this. For example, $(C[0,1], d)$ is not a complete metric space when d is the integral metric given by

$$d(f,g) = \int_0^1 |g(x) - f(x)|\, dx.$$

Example 2.9.

It seems difficult to identify a complete metric space that contains $(C[0,1], d)$. We seek a method that allows us to *construct* a complete metric space containing any given metric space as a dense subspace. This is not an easy thing to do, because we shall have to construct our space out of our (presumably incomplete) starting space.

In Section 2, we introduced the notion of *isometry* and observed that isometric metric spaces are metrically the same: any properties that depend only on the metric are common to both spaces. In particular, we saw in Theorem 2.12 that completeness is an *isometric invariant*: if (X, d_X) and (Y, d_Y) are isometric, then one is complete if, and only if, the other is. Hence it is enough for us to construct a complete metric space (X^*, d^*) that contains a dense subspace that is isometric to (X, d). Such a metric space (X^*, d^*) is called a *completion* of (X, d); see Figure 4.1.

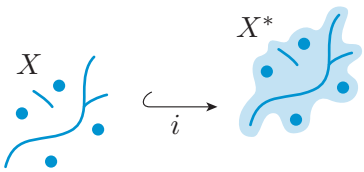

$$X^* = \mathrm{Cl}_{(X^*,d^*)}(i(X))$$

Figure 4.1

Definition 4.1 *Completion of a metric space*

Let (X, d) be a metric space. A **completion** of X is a complete metric space (X^*, d^*) that contains a dense subspace isometric to (X, d).

Remarks

1. If (X^*, d^*) and (Y^*, e^*) are both completions of (X, d), then they are isometric – thus, up to isometry, completions are unique. Consequently, we usually talk about *the* completion of a space.

We shall not prove this.

2. Note that $([0,1], d^{(1)})$ is the completion of $((0,1], d^{(1)})$ and $(\mathbb{R}, d^{(1)})$ is the completion of $(\mathbb{Q}, d^{(1)})$.

Thus our task for the remainder of this chapter is to find a procedure that allows us to construct the completion of an arbitrary metric space. This is a conceptually difficult process and you should not be surprised if you find the remainder of this chapter heavy going. Your goal should be to understand the general approach and you should not be concerned if you find the details rather technical.

4.2 Constructing the completion of a space

Given a metric space (X, d), we aim to find a complete metric space (X^*, d^*) that contains a dense subspace isometric to (X, d). The key is to find a way of constructing a new metric space out of X that is 'bigger' than X and includes the limits of every Cauchy sequence in X.

The fundamental idea is to use the Cauchy sequences in X as 'points' in the new space X^*, so that the Cauchy sequence $(a, a, a, \ldots)$ represents the point $a \in X$, and a Cauchy sequence that does not converge in X is the point in X^* that will represent the limit of the sequence; for example, $(1, 1.4, 1.41, 1.414, 1.4142, \ldots)$, which is a Cauchy sequence in $(\mathbb{Q}, d^{(1)})$ that does not converge in $\mathbb{Q}$, should represent the point $\sqrt{2}$ in the completion.

Thus let X_C denote the collection of all the d-Cauchy sequences in X. So $\mathbf{a} = (a_n) = (a_1, a_2, a_3, a_4, \ldots) \in X_C$ if, and only if, for each $\varepsilon > 0$, there is an $N \in \mathbb{N}$ such that

$$d(a_n, a_m) < \varepsilon \text{ for each } n, m > N.$$

Exercise 4.2

When $(X, d) = (\mathbb{Q}, d^{(1)})$, which of the following are points in X_C?

(a) $(0, 0, 0, 0, \ldots)$ (b) $(1/n)_{n\in\mathbb{N}}$ (c) $(n)_{n\in\mathbb{N}}$

Unfortunately, the idea of defining the completion to be the set of all Cauchy sequences in X (with an appropriate metric) does not quite work. The problem is that two distinct Cauchy sequences may have the same limit, and so they should represent the same point. For example, the sequences $(0, 0, 0, \ldots)$ and $(1/n)_{n\in\mathbb{N}}$ should both represent the point 0.

We can perhaps better understand the problem by drawing an analogy with the problem of identifying $\{(p, q) : p, q \in \mathbb{Z}, q > 0\}$ with the set of rational numbers. We would like to say that if $p, q \in \mathbb{Z}$ with $q > 0$, then the pair (p, q) represents the rational number p/q. However, this representation is not unique, since the pairs $(1, 2)$ and $(2, 4)$ both represent the rational number $\frac{1}{2}$. Instead, we define an equivalence relation on $\{(p, q) : p, q \in \mathbb{Z}, q \neq 0\}$ by defining $(p, q) \sim (p', q')$ if $pq' = p'q$, and we then represent a rational number by the corresponding equivalence class.

We do the same here, and partition X_C into disjoint equivalence classes with each class representing a different point of our completion. A class will consist of all the Cauchy sequences with the same 'limit' – or at least it would, if we had a limit to refer to! But the whole point of this construction is that we do not have *a priori* limits for all Cauchy sequences when (X, d) is incomplete.

So the next stage of the construction is to determine an equivalence relation on X_C that splits X_C into equivalence classes of Cauchy sequences that 'should' have the same limit. The following lemma helps us to do this.

Lemma 4.2

Let (X, d) be a metric space. If (a_n) and (b_n) are sequences in X that converge to the same limit, then $d(a_n, b_n) \to 0$ as $n \to \infty$.

Proof Suppose that (a_n) and (b_n) are sequences in X that converge to the same limit $a \in X$, say, and let $\varepsilon > 0$ be given. Then

$$\text{there is an } N \in \mathbb{N} \text{ such that } d(a_n, a) < \tfrac{1}{2}\varepsilon \text{ for each } n > N,$$

and

$$\text{there is an } M \in \mathbb{N} \text{ such that } d(b_n, a) < \tfrac{1}{2}\varepsilon \text{ for each } n > M.$$

Hence, whenever $n > \max\{N, M\}$, we have $d(a_n, a) < \frac{1}{2}\varepsilon$ and $d(b_n, a) < \frac{1}{2}\varepsilon$, and, by the Triangle Inequality for d,

$$d(a_n, b_n) \leq d(a_n, a) + d(a, b_n) < \tfrac{1}{2}\varepsilon + \tfrac{1}{2}\varepsilon = \varepsilon.$$

Thus $d(a_n, b_n) \to 0$ as $n \to \infty$. ■

We use this lemma to define a relation on X_C: if $\mathbf{a} = (a_n) \in X_C$ and $\mathbf{b} = (b_n) \in X_C$, then we say that $\mathbf{a}$ is equivalent to $\mathbf{b}$ (denoted $\mathbf{a} \sim \mathbf{b}$) if, and only if, $d(a_n, b_n) \to 0$ as $n \to \infty$.

Exercise 4.3

Show that $\sim$ is an equivalence relation on X_C.

We now define X^* to be the collection of equivalence classes of $(X_C, \sim)$.

We still need to:

1. define a metric d^* on X^*
2. show that (X, d) is isometric to a dense subset of (X^*, d^*)
3. show that (X^*, d^*) is complete.

Defining a metric on X^*

Let $\mathbf{a} = (a_1, a_2, a_3, \ldots)$ be a Cauchy sequence in X, and let $[\mathbf{a}]$ denote the equivalence class of $\mathbf{a}$. Thus,

$$[\mathbf{a}] = \{\mathbf{a}' = (a'_1, a'_2, a'_3, \ldots) \in X_C : d(a_n, a'_n) \to 0 \text{ as } n \to \infty\}.$$

Given two points $[\mathbf{a}], [\mathbf{b}] \in X^*$, we define our metric d^* as follows:

$$d^*([\mathbf{a}], [\mathbf{b}]) = \lim_{n\to\infty} d(a_n, b_n), \tag{1}$$

where $(a_n) \in [\mathbf{a}]$ and $(b_n) \in [\mathbf{b}]$.

First, we must verify that d^* is well defined – that is, the limit in (1) exists, and if $\mathbf{a}' \in [\mathbf{a}]$ and $\mathbf{b}' \in [\mathbf{b}]$ are different choices of sequences from the equivalence classes, then

$$\lim_{n\to\infty} d(a'_n, b'_n) = \lim_{n\to\infty} d(a_n, b_n).$$

For otherwise, d^* would depend on which representatives of the equivalence class are used to calculate the limit, and so would not be well defined.

Lemma 4.3

Let $\mathbf{a} = (a_n) \in X_C$ and $\mathbf{b} = (b_n) \in X_C$. Then

1. $\lim_{n\to\infty} d(a_n, b_n)$ exists
2. if $\mathbf{a}' = (a'_n) \in [\mathbf{a}]$, $\mathbf{b}' = (b'_n) \in [\mathbf{b}]$, then
$$\lim_{n\to\infty} d(a'_n, b'_n) = \lim_{n\to\infty} d(a_n, b_n).$$

Proof

1. To show that $\lim_{n\to\infty} d(a_n, b_n)$ exists, we prove that $(d(a_n, b_n))$ is a Cauchy sequence in $\mathbb{R}$ for the Euclidean metric. For $n, m \in \mathbb{N}$, we use the Triangle Inequality for d, followed by the Reverse Triangle Inequality, as follows:
$$\begin{aligned} 0 &\le |d(a_n, b_n) - d(a_m, b_m)| \\ &= |(d(a_n, b_n) - d(a_m, b_n)) + (d(a_m, b_n) - d(a_m, b_m))| \\ &\le |d(a_n, b_n) - d(a_m, b_n)| + |d(a_m, b_n) - d(a_m, b_m)| \\ &\le d(a_n, a_m) + d(b_n, b_m). \end{aligned}$$
But (a_n) and (b_n) are both Cauchy sequences, and so $(d(a_n, b_n))$ is also a Cauchy sequence. Since $(\mathbb{R}, d^{(1)})$ is complete, $(d(a_n, b_n))$ converges.
2. Now suppose that $\mathbf{a}' \in [\mathbf{a}]$ and $\mathbf{b}' \in [\mathbf{b}]$.

   To prove that d^* is defined independently of the chosen representatives of the equivalence class, we show that
   $$|d(a'_n, b'_n) - d(a_n, b_n)| \to 0 \text{ as } n \to \infty.$$
   Since $\mathbf{a}' \in [\mathbf{a}]$ and $\mathbf{b}' \in [\mathbf{b}]$, we have $d(a'_n, a_n) \to 0$ and $d(b'_n, b_n) \to 0$.

Thus we use the Triangle Inequality, followed by the Reverse Triangle Inequality to estimate

$$\begin{aligned}|d(a'_n, b'_n) - d(a_n, b_n)| &= |d(a'_n, b'_n) - d(a'_n, b_n) + d(a'_n, b_n) - d(a_n, b_n)| \\ &\leq |d(a'_n, b'_n) - d(a'_n, b_n)| + |d(a'_n, b_n) - d(a_n, b_n)| \\ &\leq d(b'_n, b_n) + d(a'_n, a_n).\end{aligned}$$

But $d(b'_n, b_n) \to 0$ and $d(a'_n, a_n) \to 0$ as $n \to \infty$. Hence $|d(a'_n, b'_n) - d(a_n, b_n)| \to 0$ as $n \to \infty$ and the result follows.

■

Now we know that d^* is well defined, we can show that it defines a metric on X^*.

Theorem 4.4

(X^*, d^*) is a metric space.

Exercise 4.4

Prove Theorem 4.4.

Show that (X, d) is isometric to a dense subset of (X^*, d^*)

We start by considering which elements of X^* can be identified with the points of X. If $x \in X$, then $\mathbf{x} = (x, x, x, \ldots) \in X_C$ is a Cauchy sequence in X that converges to x. This gives us an embedding $i\colon X \to X^*$ where

$$i(x) = [(x, x, x, \ldots)] = [\mathbf{x}].$$

Thus $i(x)$ is the equivalence class consisting of all Cauchy sequences (a_n) in X for which $d(a_n, x) \to 0$ as $n \to \infty$ – that is, $i(x)$ consists of all sequences that converge to x.

Clearly for each $x, y \in X$,

$$d^*(i(x), i(y)) = \lim_{n\to\infty} d(x, y) = d(x, y),$$

and so $i(x) = i(y)$ if, and only if, $d(x, y) = 0$ which (by (M1)) holds if, and only if, $x = y$. Hence i is one–one; this implies that $i(X) = \{i(x) : x \in X\}$ is an isometric copy of X in X^*, and so (X, d) is isometric to $(i(X), d^*_{i(X)})$.

In fact, $i(X)$ is d^*-dense in X^*, as we now prove.

Lemma 4.5

The set $i(X)$ is a d^*-dense subset of (X^*, d^*).

Proof It is enough to show that for each $a^* \in X^*$ and each $\varepsilon > 0$, there is an $x \in X$ with $d^*(a^*, i(x)) < \varepsilon$.

See Book D, Chapter 16, Theorem 3.8.

So let $a^* \in X^*$, and let $\varepsilon > 0$ be given.

Since $a^* \in X^*$, there is a Cauchy sequence $\mathbf{a} = (a_n)$ in X such that $a^* = [\mathbf{a}]$. Choose an $N \in \mathbb{N}$ such that $d(a_n, a_m) < \frac{1}{2}\varepsilon$ for each $m, n > N$. Then for $x = a_{N+1}$,

$$d^*(a^*, i(x)) = d^*(a^*, i(a_{N+1})) = \lim_{n\to\infty} d(a_n, a_{N+1}).$$

But for $n > N$, $d(a_n, a_{N+1}) < \frac{1}{2}\varepsilon$, and so

$$d^*(a^*, i(x)) \leq \tfrac{1}{2}\varepsilon < \varepsilon.$$

Thus $i(X)$ is d^*-dense in X^*. ■

Showing that (X^*, d^*) is complete

Theorem 4.6

The metric space (X^*, d^*) is complete.

Proof We need to show that each d^*-Cauchy sequence in X^* converges to an element of X^*.

So suppose that $(A_m)_{m\in\mathbb{N}}$ is a d^*-Cauchy sequence in X^*. Each term A_m of the sequence is an equivalence class of Cauchy sequences from X. Hence for each $m \in \mathbb{N}$, we can choose a representative of the equivalence class, $\mathbf{a}_m = (a_{1,m}, a_{2,m}, a_{3,m}, \ldots)$ say, so that $A_m = [\mathbf{a}_m]$. Moreover, for each $m \in \mathbb{N}$, $(a_{n,m})_{n\in\mathbb{N}}$ is itself a d-Cauchy sequence in X.

Thus $(A_m)_{m\in\mathbb{N}}$ is a Cauchy sequence consisting of sets of Cauchy sequences!

We first identify a candidate for the limit of the sequence $([\mathbf{a}_m])$. To do this, we use Lemma 4.5 to find for each $m \in \mathbb{N}$, a point $x_m \in X$ such that

$$d^*(A_m, i(x_m)) = d^*([\mathbf{a}_m], i(x_m)) < \frac{1}{m}.$$

If we show that

1. $(x_m)_{m\in\mathbb{N}}$ is a d-Cauchy sequence in X,
2. $d^*(A_m, [(x_1, x_2, x_3, \ldots)]) = d^*([\mathbf{a}_m], [(x_1, x_2, x_3, \ldots)]) \to 0$ as $m \to \infty$,

then we can deduce that (A_m) is d^*-convergent (to $[(x_1, x_2, x_3, \ldots)]$) in X^*, as required.

$(x_m)_{m\in\mathbb{N}}$ is a d-Cauchy sequence in X

Let $\varepsilon > 0$. For each $m, n \in \mathbb{N}$ and each $j \in \mathbb{N}$ we can use the Triangle Inequality for d to estimate

$$d(x_m, x_n) \leq d(x_m, a_{j,m}) + d(a_{j,m}, a_{j,n}) + d(a_{j,n}, x_n).$$

But

$$\lim_{j\to\infty} d(a_{j,m}, x_m) = d^*([\mathbf{a}_m], i(x_m)) = d^*(A_m, i(x_m)) < \frac{1}{m}$$

and

$$\lim_{j\to\infty} d(a_{j,n}, x_n) = d^*([\mathbf{a}_n], i(x_n)) = d^*(A_n, i(x_n)) < \frac{1}{n}.$$

Moreover, $(A_m) = ([\mathbf{a}_m])_{m\in\mathbb{N}}$ is a Cauchy sequence in X^*, and so there is an $N \in \mathbb{N}$ such that for each $n, m > N$,

$$d^*(A_m, A_n) = d^*([\mathbf{a}_m], [\mathbf{a}_n]) = \lim_{j\to\infty} d(a_{j,m}, a_{j,n}) < \tfrac{1}{3}\varepsilon.$$

So if we choose $n, m > \max\{N, 3/\varepsilon\}$, and choose $j \in \mathbb{N}$ so that $d(a_{j,m}, a_{j,n}) < \frac{1}{3}\varepsilon$, then

$$\begin{aligned} d(x_m, x_n) &\leq d(x_m, a_{j,m}) + d(a_{j,m}, a_{j,n}) + d(a_{j,n}, x_n) \\ &< \frac{1}{m} + \frac{1}{3}\varepsilon + \frac{1}{n} \\ &< \frac{1}{3}\varepsilon + \frac{1}{3}\varepsilon + \frac{1}{3}\varepsilon = \varepsilon. \end{aligned}$$

Hence, $(x_m)_{m\in\mathbb{N}}$ is a Cauchy sequence in X.

$d^*(A_m, [(x_1, x_2, x_3, \ldots)]) = d^*([\mathbf{a}_m], [(x_1, x_2, x_3, \ldots)]) \to 0$ *as* $m \to \infty$

Let $\varepsilon > 0$ be given. Then for each $m \in \mathbb{N}$, we use the Triangle Inequality for d^* to estimate

$$\begin{aligned} d^*([\mathbf{a}_m], [(x_1, x_2, x_3, \ldots)]) &\leq d^*([\mathbf{a}_m], i(x_m)) + d^*(i(x_m), [(x_1, x_2, x_3, \ldots)]) \\ &< \frac{1}{m} + \lim_{j\to\infty} d(x_m, x_j). \end{aligned}$$

But $(x_j)_{j\in\mathbb{N}}$ is a d-Cauchy sequence, and so we can find an $N \in \mathbb{N}$ such that $d(x_m, x_j) < \frac{1}{2}\varepsilon$ for each $j, m > N$. Hence if we choose $m > \max\{N, 2/\varepsilon\}$, then

$$d^*([\mathbf{a}_m], [(x_1, x_2, x_3, \ldots)]) < \tfrac{1}{2}\varepsilon + \tfrac{1}{2}\varepsilon = \varepsilon,$$

and so $A_m = [\mathbf{a}_m] \xrightarrow{d^*} [(x_1, x_2, x_3, \ldots)]$ as $m \to \infty$, as required. ■

Our results are summarised in the following theorem.

Theorem 4.7

Any metric space (X, d) has a completion (X^*, d^*).

We have shown that any metric space can be embedded as a dense subset of some complete metric space. However, the abstract construction that we have given is of limited utility in practice – it is useful for showing that the completion exists, but is hard to understand and make use of for concrete problems. It does not give us a clear understanding of the structure of the completion for any particular space: this must be obtained by special methods. In fact, only a few completions are known in detail, such as $\mathbb{Q}^* = \mathbb{R}$. However, if such a representation is found, it is unique up to isometry.

We omit the proof of this.

Returning briefly to $(C[0,1], d)$, where d is the integral metric, Theorem 4.7 tells us that there is a complete metric space (X^*, d^*) that contains $(C[0,1], d)$ as a dense subspace. In fact, the completion of $(C[0,1], d)$ is the space known as $L^1[0,1]$, the set of *Lebesgue integrable functions* on the interval $[0,1]$, and the metric is based on an extension of the definition of the integral for continuous functions known as the Lebesgue integral. This space plays a fundamental role in the theory of modern analysis, but study of it is beyond the scope of this course.

$d(f,g) = \int_0^1 |f(x) - g(x)|\, dx.$

Remark

The Lebesgue integral is a generalisation of the Riemann integral – it agrees with the Riemann integral for continuous functions, but is also defined for many functions that are not continuous, such as the function $f\colon [0,1] \to \mathbb{R}$ given by $f(x) = 1$ if X is rational, and 0 otherwise – the Lebesgue integral of f is 0.

Learning outcomes

We list those topics on which we may set assessment questions to test your understanding of the chapter.

After working through this chapter you should be able to:

(a) Test whether given sequences in a metric space are Cauchy.

(b) Determine when given subsets of standard metric spaces are complete.

(c) Recognise standard complete metric spaces.

(d) Use the Contraction Mapping Theorem to determine whether certain maps have a unique fixed point, and to find approximations to the fixed point that are within a specified distance.

(e) Understand the approach to constructing the completion of a given metric space.

Solutions to exercises

Solution to Exercise 1.1

(a) If (o_n) is the odd subsequence, then

$$\begin{aligned} o_n = b_{2n-1} &= -1 + \frac{1}{2} - \frac{1}{3} + \cdots + \frac{(-1)^{2n-1}}{2n-1} \\ &= -1 + \frac{1}{2} - \frac{1}{3} + \cdots - \frac{1}{2n-1} \end{aligned}$$

for $n \in \mathbb{N}$ and so

$$o_{n+1} - o_n = -\frac{1}{2n+1} + \frac{1}{2n} = \frac{1}{4n^2+2n} > 0 \text{ for } n \in \mathbb{N}.$$

Hence (o_n) is monotonic increasing. Moreover, for each $n \in \mathbb{N}$,

$$o_n = \left(-1 + \frac{1}{2}\right) + \left(-\frac{1}{3} + \frac{1}{4}\right) + \cdots + \frac{-1}{2n-1} < 0$$

and so, by the Monotone Convergence Theorem (o_n) is a convergent sequence.

Similarly, if (e_n) is the even subsequence, then

$$\begin{aligned} e_n = b_{2n} &= -1 + \frac{1}{2} - \frac{1}{3} + \cdots + \frac{(-1)^{2n}}{2n} \\ &= -1 + \frac{1}{2} - \frac{1}{3} + \cdots + \frac{1}{2n} \end{aligned}$$

for $n \in \mathbb{N}$ and so

$$\begin{aligned} e_{n+1} - e_n &= \frac{-1}{2n+1} + \frac{1}{2n+2} \\ &= \frac{-1}{2(n+1)(2n+1)} < 0 \text{ for } n \in \mathbb{N}. \end{aligned}$$

Hence (e_n) is monotonic decreasing. Moreover, for each $n \in \mathbb{N}$,

$$\begin{aligned} e_n &= \left(-1 + \frac{1}{2}\right) + \left(-\frac{1}{3} + \frac{1}{4}\right) + \cdots + \left(\frac{-1}{2n-1} + \frac{1}{2n}\right) \\ &= -\frac{1}{2} - \frac{1}{12} - \cdots - \frac{1}{2n(2n-1)} \geq -\frac{1}{1} - \frac{1}{3^2} - \cdots - \frac{1}{4n^2} \\ &\geq -a_{2n} \geq -2 \end{aligned}$$

and so, by the Monotone Convergence Theorem (e_n) is a convergent sequence.

(b) For each $n \in \mathbb{N}$, we find

$$\begin{aligned} d_n &= \left(-1 + \frac{1}{2} - \frac{1}{3} + \cdots + \frac{(-1)^{2n-1}}{2n-1}\right) - \left(-1 + \frac{1}{2} - \frac{1}{3} + \cdots + \frac{(-1)^{2n}}{2n}\right) \\ &= -\frac{1}{2n}, \end{aligned}$$

which is a null sequence.

(c) Since d_n is null, it follows that both (o_n) and (e_n) converge to the same limit. In other words, both the odd and even subsequences of (b_n) converge to the same limit, which implies that (b_n) is also convergent (and has the same limit).

Solution to Exercise 1.2

(a) If $n, m > N$, then

$$\left|\frac{1}{n} - \frac{1}{m}\right| \leq \frac{1}{n} + \frac{1}{m} < \frac{2}{N}.$$

Now let $\varepsilon > 0$. Choose $N \in \mathbb{N}$ so that $N > \frac{2}{\varepsilon}$. If $n, m > N$, then

$$\left|\frac{1}{n} - \frac{1}{m}\right| < \frac{2}{N} < \varepsilon.$$

So $(1/n)$ is a Cauchy sequence.

Alternatively, $(1/n)$ is d-convergent with limit 0; so, by Lemma 1.2, it is a d-Cauchy sequence.

(b) Let $\varepsilon = 1$ and let $N \in \mathbb{N}$. Then

$$|(N+2) - (N+1)| = 1 = \varepsilon$$

and so $(n)_{n\in\mathbb{N}}$ is not a Cauchy sequence. (We could have used a similar argument for any value of $\varepsilon > 0$.)

Solution to Exercise 1.3

By Worked Exercise 3.5 of Chapter 14, eventually constant sequences are convergent; so, by Lemma 1.2, we deduce that they are also Cauchy.

Solution to Exercise 1.4

(a) To show that (c_n) is a Cauchy sequence, we note that

$$c_{n+1} - c_n = \frac{(-1)^{n+1}}{2^{n+1}}.$$

Hence if $m = n + i$ for some $i \in \mathbb{N}$, then

$$\begin{aligned} |c_{n+i} - c_n| &= |(c_{n+i} - c_{n+i-1}) + (c_{n+i-1} - c_{n+i-2}) + \cdots + (c_{n+1} - c_n)| \\ &\leq \frac{1}{2^{n+i}} + \frac{1}{2^{n+i-1}} + \cdots + \frac{1}{2^{n+1}} \\ &\leq \frac{1}{2^n}. \end{aligned}$$

Thus for each $\varepsilon > 0$, if we choose $N \in \mathbb{N}$ so that $2^{-N} < \varepsilon$, then for $n, m > N$, we have

$$|a_n - a_m| < \varepsilon$$

and so (c_n) is Cauchy as required.

(b) If (e_n) denotes the subsequence of even terms, so $e_n = c_{2n}$, then

$$e_{n+1} - e_n = -\frac{1}{2^{2n+2}} < 0,$$

and so (e_n) is a monotonic decreasing sequence. Moreover, $e_n \geq -\sum_{i=1}^{2n} \frac{1}{2^i} \geq -2$ and so it is bounded below. Hence by the Monotone Convergence Theorem (e_n) is a convergent sequence and we have found a convergent subsequence of (c_n).

We deduce from Lemma 1.6 that (c_n) is convergent.

Solution to Exercise 2.1

In Exercise 1.2 you saw that the sequence $(1/n)$ is a $d^{(1)}$-Cauchy sequence in $\mathbb{R}$. Since each term of the sequence lies in $(0, 1]$, it follows that the sequence is also Cauchy for the induced metric on $(0, 1]$. However, although each term of this sequence lies in $(0, 1]$, its limit 0 is not a member of $(0, 1]$ and so the sequence does not converge in $(0, 1]$. Thus, $(0, 1]$ with the induced Euclidean metric is incomplete.

Solution to Exercise 2.2

Certainly it is the case that convergent sequences are Cauchy (in any metric space) and so it only remains to show that Cauchy sequences are convergent. So suppose that (a_n) is a Cauchy sequence in (X, d_0). Then, in particular, there is $N \in \mathbb{N}$ such that for each $n, m > N$, $d_0(a_n, a_m) < 1$. But, since d_0 is the discrete metric for X, this means that for each $n, m > N$, $a_n = a_m$. In other words, the sequence (a_n) is *eventually constant*. But by Worked Exercise 3.5 of Chapter 14, eventually constant sequences are convergent (whatever the choice of metric) and so (X, d_0) is complete as required.

Solution to Exercise 2.3

The set $[0, 1]$ is the only one of the sets given that is closed for the Euclidean metric, and so it is the only set A for which $(A, d_A^{(1)})$ is complete. By considering the sequence $(1 - \frac{1}{n})$ for $[0, 1)$ and the sequence $(\frac{1}{n+1})$ for $(0, 1)$, it is possible to verify directly that neither of these sets is complete for the metric induced from the Euclidean metric.

Solution to Exercise 2.4

(a) The continuous image $f(X)$ of a compact space X is itself compact, by Corollary 5.13 of Chapter 22, and so, since $f(X) = Y$, (Y, e) is compact. Hence, by Theorem 2.3, it is complete.

(b) If f is not onto, then $f(X)$ is a compact *proper* subset of Y and hence complete but without further information, we cannot determine whether (Y, e) is itself complete.

Solution to Exercise 2.5

See Exercise 4.7 of Chapter 22.

Since finite sets are compact, (X, d) is a compact metric space and the result follows immediately from Theorem 2.3.

Solution to Exercise 2.6

We know that the closed unit ball $A = B_{d^{(n)}}[\mathbf{0}, 1]$ is $d^{(n)}$-closed (see Proposition 1.3 of Chapter 16) and that $(\mathbb{R}^n, d^{(n)})$ is complete (see Corollary 2.5). Hence by Proposition 2.2 we conclude that $(A, d_A^{(n)})$ is complete.

Solution to Exercise 2.7

Let $(\mathbf{x}_n)$ be a $d_{\mathbf{C}}$-Cauchy sequence in $\mathbf{C}$ and let $k \in \mathbb{N}$ be given.

Then for each $\varepsilon > 0$, there is $N \in \mathbb{N}$ such that for $n, m > N$, $d_{\mathbf{C}}(\mathbf{x}_n, \mathbf{x}_m) < \varepsilon$. In particular, if we let $\varepsilon = 2^{-k}$, then there is $N_k \in \mathbb{N}$ such that for $n, m > N_k$, $d_{\mathbf{C}}(\mathbf{x}_n, \mathbf{x}_m) < 2^{-k}$. In particular, we must have $\mathbf{x}_n$ and $\mathbf{x}_m$ agreeing for their first k terms.

Solution to Exercise 2.8

(a) We must find an isometry from $[0, 1]$ to $[1, 2]$. Define $f\colon [0, 1] \to [1, 2]$ by $f(x) = x + 1$. It is easy to see that f is one–one and onto, and it remains only to verify that f preserves distances.

So let $x, y \in [0, 1]$; then

$$\begin{aligned} d_{[1,2]}(f(x), f(y)) &= |(x + 1) - (y + 1)| \\ &= |x - y| = d_{[0,1]}(x, y), \end{aligned}$$

and so distances are preserved by f.

Hence f is an isometry and $([0, 1], d_{[0,1]})$ and $([1, 2], d_{[1,2]})$ are isometric.

(b) To show that two spaces are not isometric, it is enough to show that there are points in one space that are at a distance not realised in the other space. In this case, the distance between the points $0, 2 \in [0, 2]$ is 2, whereas the distance between any two points in $[0, 1]$ is at most 1. Hence the two spaces are not isometric.

Solution to Exercise 2.9

Suppose that (a_n) is a sequence in X that is d_X-convergent to a. Then, since f is an isometry,

$$d_Y(f(a_n), f(a)) = d_X(a_n, a) \to 0 \text{ as } n \to \infty,$$

and so $(f(a_n))$ is a sequence in Y that is d_Y-convergent to $f(a)$.

Solution to Exercise 3.1

The point X is a fixed point of g if, and only if,

$$g(X) = X,$$

that is, if, and only if,

$$X - f(X) = X.$$

This is equivalent to $f(X) = 0$. Thus, X is a fixed point of g if, and only if, X is a zero of f.

Solution to Exercise 3.2

The point X is a fixed point of g_r if, and only if,

$$g_r(X) = X,$$

that is, if, and only if,

$$X - rf(X) = X.$$

Since $r \neq 0$, this is equivalent to $f(X) = 0$. Thus, X is a fixed point of g_r if, and only if, X is a zero of f.

Solution to Exercise 3.3

Integrating both sides of the differential equation with respect to s gives

$$\int_0^t \frac{df}{ds}\, ds = \int_0^t \tfrac{1}{2}\cos f(s)\, ds$$

and so

$$f(t) - f(0) = \int_0^t \tfrac{1}{2}\cos f(s)\, ds, \text{ for } 0 \leq t \leq 1.$$

Since $f(0) = 0$, this gives the integral equation

$$f(t) = \int_0^t \tfrac{1}{2}\cos f(s)\, ds, \text{ for } 0 \leq t \leq 1.$$

Thus, if we define a map T by

$$T(f)(t) = \int_0^t \tfrac{1}{2}\cos f(s)\, ds, \text{ for } 0 \leq t \leq 1,$$

then any solution of the differential equation is also a fixed point of this map T.

We now show that T is a map from $C[0,1]$ to itself. Since $s \mapsto (\frac{1}{2}\cos \circ f)(s)$ is a composite of continuous functions, it is continuous. Also, the map F defined by

$$F(g)(t) = \int_0^t g(s)\, ds, \text{ for } 0 \leq t \leq 1,$$

is a map from $C[0,1]$ to itself (as shown in Book D, Chapter 15, Worked Exercise 5.7) and so T is a map from $C[0,1]$ to itself. All solutions of the differential equation correspond to fixed points of the integral equation $T(f) = f$.

Solution to Exercise 3.4

Let $a \in X$ and let $c = T(a)$. We show that $T(x) = c$, for each $x \in X$.

Let us estimate $d(T(x), T(a))$. Since T has contraction ratio 0,

$$0 \leq d(T(x), c) = d(T(x), T(a)) \leq 0 \times d(x, a) = 0,$$

for each $x \in X$. Hence $d(T(x), c) = 0$ and so, since d is a metric, $T(x) = c$. In other words, T is a constant function.

Solution to Exercise 3.5

For $f, g \in C[0,1]$ and $t \in [0,1]$, we have

$$\begin{aligned}
|T(f)(t) - T(g)(t)| &= \left| \int_0^t \tfrac{1}{2}(\cos f(s) - \cos g(s))\, ds \right| \\
&\leq \tfrac{1}{2} \int_0^t |\cos f(s) - \cos g(s)|\, ds \\
&\leq \tfrac{1}{2} \int_0^t |f(s) - g(s)|\, ds, \text{ using the hint} \\
&\leq \tfrac{1}{2} \int_0^t d_{\max}(f,g)\, ds \\
&= \tfrac{1}{2} t d_{\max}(f,g) \leq \tfrac{1}{2} d_{\max}(f,g), \text{ since } 0 \leq t \leq 1.
\end{aligned}$$

Hence,

$$d_{\max}(T(f), T(g)) \leq \tfrac{1}{2} d_{\max}(f,g),$$

and so T is a contraction mapping for $d_{\max}$ with contraction ratio $\frac{1}{2}$.

Solution to Exercise 3.6

(a) It is enough to show that T is a Lipschitz mapping on $(0,1]$ with Lipschitz constant $\frac{1}{2}$.

If $x, y \in (0,1]$, then

$$|T(x) - T(y)| = |\tfrac{1}{2}x - \tfrac{1}{2}y| = \tfrac{1}{2}|x - y|.$$

Hence T is a Lipschitz mapping on $(0,1]$ with Lipschitz constant $\frac{1}{2}$.

(b) If $T(x) = x$, then $\frac{1}{2}x = x$; that is, $\frac{1}{2}x = 0$. So the only possible fixed point of T is the point 0, which does not belong to the domain of T, $(0,1]$. Thus T has no fixed points.

Solution to Exercise 3.7

(a) For each $x, y \in [1, \infty)$,

$$\begin{aligned}
|T(x) - T(y)| &= \left| \left(x + \frac{1}{x}\right) - \left(y + \frac{1}{y}\right) \right| \\
&= \left| x - y + \frac{y - x}{xy} \right| \\
&= \left| (x - y)\left(1 - \frac{1}{xy}\right) \right| \\
&= |x - y| \times \left| 1 - \frac{1}{xy} \right| \\
&< |x - y|.
\end{aligned}$$

(b) If T has a fixed point X, then $x = T(x)$. This means that $x = x + 1/x$, and so $1/x = 0$, which is impossible. Thus T has no fixed points.

(c) This does not contradict the Contraction Mapping Theorem, because T is not a contraction mapping with contraction ratio strictly less than 1.

Solution to Exercise 3.8

By Corollary 3.8,

$$|x_T - T^n(0)| \leq \frac{\left(\frac{1}{2}\right)^n}{1-\frac{1}{2}}|0 - T(0)| = \left(\frac{1}{2}\right)^n.$$

This places x_T in an interval of length $(1/2)^{n-1}$. Now

$$((1/2)^{n-1}) = (1, \tfrac{1}{2}, \tfrac{1}{4}, \tfrac{1}{8}, \tfrac{1}{16}, \tfrac{1}{32}, \cdots).$$

The first term in this sequence which is less than 0.1 is the term corresponding to $n = 5$. We have

$$|x_T - T^5(0)| \leq \left(\frac{1}{2}\right)^5 = \frac{1}{32}.$$

Since $T^5(0) \in (0.4502, 0.4503)$ and $\frac{1}{32} = 0.03125$,

$$x_T \in [T^5(0) - \tfrac{1}{32}, T^5(0) + \tfrac{1}{32}] \subseteq [0.41895, 0.48155] \subseteq (0.41, 0.49).$$

This places x_T in an interval of length less than 0.1.

Solution to Exercise 3.9

We follow the method used in the text before the theorem. Since $0 < m \leq f'(x) \leq M$ on $[a, b]$, we find for each $x \in [a, b]$,

$$1 - Mr \leq g_r'(x) = 1 - rf'(x) \leq 1 - mr.$$

Hence, if $0 < r \leq \frac{1}{M}$, then $1 - Mr \geq 0$, and so g_r is non-decreasing on $[a, b]$. In this case, for $x \in [a, b]$,

$$g_r(x) \geq g_r(a) = a - rf(a) > a$$

and

$$g_r(x) \leq g_r(b) = b - rf(b) < b.$$

Hence, for $0 < r \leq \frac{1}{M}$, g_r maps $[a, b]$ into itself. Finally, we apply the Mean Value Theorem to g_r to conclude that if $0 < r \leq \frac{1}{M}$, then for $x, y \in [a, b]$,

$$\begin{aligned} |g(x) - g(y)| &\leq \max_{c \in [a,b]} |g_r'(c)||x - y| \\ &\leq (1 - mr)|x - y|, \end{aligned}$$

and so g_r is a contraction mapping (for the Euclidean metric) on $[a, b]$ with contraction ratio $1 - mr$.

Solution to Exercise 3.10

For each $t \in [0,1]$,

$$\begin{aligned} T(f)(t) &= 1 + \int_0^t sf(s)\,ds \\ &= 1 + \int_0^t se^{\frac{1}{2}s^2}\,ds \\ &= 1 + \left[e^{\frac{1}{2}s^2}\right]_0^t \\ &= 1 + \left(e^{\frac{1}{2}t^2} - 1\right) = f(t). \end{aligned}$$

Hence $T(f) = f$, as required.

Solution to Exercise 3.11

(a) By the result of Exercise 3.3, any solution to the differential equation must also be a fixed point of the mapping $T\colon C[0,1] \to C[0,1]$, given by

$$T(f)(t) = \int_0^t \tfrac{1}{2}\cos f(s)\,ds \text{ for each } 0 \leq t \leq 1.$$

In Exercise 3.5, we saw that this is a contraction mapping (for $d_{\max}$) with contraction ratio $\frac{1}{2}$. Hence by the Contraction Mapping Theorem, there is a unique function $f \in C[0,1]$ that is a fixed point of T. Since $(s,y) \mapsto \frac{1}{2}\cos y$ is a $(d^{(2)}, d^{(1)})$-continuous function, it follows from Theorem 3.10 that f is also the unique solution of the original differential equation.

(b) Taking $f_0(s) = 0$ as our initial guess for the fixed point f, we find that

$$T(f_0) = \tfrac{1}{2}t$$

and

$$T(T(f_0)) = \int_0^t \tfrac{1}{2}\cos\tfrac{1}{2}s\,ds = \sin\tfrac{1}{2}t.$$

We deduce from Corollary 3.8 that

$$\begin{aligned} d_{\max}(f, \sin\tfrac{1}{2}t) &\leq \frac{(\frac{1}{2}^2)}{1-\frac{1}{2}} d_{\max}(f_0, T(f_0)) \\ &= \frac{\frac{1}{4}}{1-\frac{1}{2}} \times \tfrac{1}{2} = \tfrac{1}{4}, \end{aligned}$$

and so we can take $g(t) = \sin\frac{1}{2}t$ as our estimate to the solution of the differential equation.

(In fact, according to Wolfram Alpha, the solution of this differential equation is given by $f(s) = 2\tan^{-1}\left(\tanh\left(\frac{s}{4}\right)\right)$.)

Solution to Exercise 4.1

In Worked Exercise 3.9 of Chapter 16, we saw that $\mathbb{Q}$ is a dense subset of $\mathbb{R}$ for the Euclidean metric, and so, since Theorem 1.8 tells us that $(\mathbb{R}, d^{(1)})$ is a complete metric space, we take $(X^*, d^*) = (\mathbb{R}, d^{(1)})$.

Solution to Exercise 4.2

Both $(0, 0, 0, 0, \ldots)$ and $(1/n)$ are Cauchy sequences of rational numbers, and so are points in X_C.

However, $(n)_{n\in\mathbb{N}}$ is not a Cauchy sequence, and so is not in X_C.

Solution to Exercise 4.3

We need to show that $\sim$ is reflexive, symmetric and transitive on X_C.

Reflexive: $\mathbf{a} \sim \mathbf{a}$. Let $\mathbf{a} = (a_1, a_2, a_3, \ldots) \in X_C$. For each $n \in \mathbb{N}$, $d(a_n, a_n) = 0$, and so $d(a_n, a_n) \to 0$. Hence $\mathbf{a} \sim \mathbf{a}$.

Symmetric: $\mathbf{a} \sim \mathbf{b}$ *if, and only if,* $\mathbf{b} \sim \mathbf{a}$. Let $\mathbf{a} = (a_1, a_2, a_3, \ldots)$ and $\mathbf{b} = (b_1, b_2, b_3, \ldots)$ belong to X_C. Since $d(a_n, b_n) = d(b_n, a_n)$ for each $n \in \mathbb{N}$, we deduce that $d(a_n, b_n) \to 0$ if, and only if, $d(b_n, a_n) \to 0$ as $n \to \infty$. Thus, $\mathbf{a} \sim \mathbf{b}$ if, and only if, $\mathbf{b} \sim \mathbf{a}$.

Transitive: if $\mathbf{a} \sim \mathbf{b}$ *and* $\mathbf{b} \sim \mathbf{c}$ *then* $\mathbf{a} \sim \mathbf{c}$. Let $\mathbf{a} = (a_1, a_2, a_3, \ldots), \mathbf{b} = (b_1, b_2, b_3, \ldots)$ and $\mathbf{c} = (c_1, c_2, c_3, \ldots)$ belong to X_C.

If $\mathbf{a} \sim \mathbf{b}$ and $\mathbf{b} \sim \mathbf{c}$, then, for each $n \in \mathbb{N}$, the Triangle Inequality for d implies that

$$0 \leq d(a_n, c_n) \leq d(a_n, b_n) + d(b_n, c_n).$$

But $d(a_n, b_n) \to 0$ and $d(b_n, c_n) \to 0$ as $n \to \infty$, and so $d(a_n, c_n) \to 0$ as $n \to \infty$. Hence, $\mathbf{a} \sim \mathbf{c}$.

Thus $\sim$ is an equivalence relation on X_C.

Solution to Exercise 4.4

We show that d^* defines a metric on X^*. In Lemma 4.3 we showed that $d^* : X^* \times X^* \to \mathbb{R}$. We now need to verify properties (M1)–(M3). Throughout the following we let $\mathbf{a} = (a_n)$, $\mathbf{b} = (b_n)$ and $\mathbf{c} = (c_n)$.

(M1) By definition, $d^*([\mathbf{a}], [\mathbf{b}]) \geq 0$, and

$$d^*([\mathbf{a}], [\mathbf{a}]) = \lim_{n\to\infty} d(a_n, a_n) = 0.$$

If $d^*([\mathbf{a}], [\mathbf{b}]) = 0$, then $d(a_n, b_n) \to 0$; this implies that $\mathbf{a} \sim \mathbf{b}$, and so $[\mathbf{a}] = [\mathbf{b}]$. Hence (M1) holds.

(M2) Since d is a metric we have, for $[\mathbf{a}], [\mathbf{b}] \in X^*$,

$$\begin{aligned} d^*([\mathbf{a}], [\mathbf{b}]) &= \lim_{n\to\infty} d(a_n, b_n) \\ &= \lim_{n\to\infty} d(b_n, a_n) = d^*([\mathbf{b}], [\mathbf{a}]), \end{aligned}$$

and so (M2) holds.

(M3) Let $[\mathbf{a}], [\mathbf{b}], [\mathbf{c}] \in X^*$. Since the Triangle Inequality for d implies

$$d(a_n, c_n) \leq d(a_n, b_n) + d(b_n, c_n),$$

and Lemma 4.3 implies that $\lim_{n\to\infty} d(a_n, c_n)$, $\lim_{n\to\infty} d(a_n, b_n)$ and $\lim_{n\to\infty} d(b_n, c_n)$ exist, it follows that

$$\lim_{n\to\infty} d(a_n, c_n) \leq \lim_{n\to\infty} d(a_n, b_n) + \lim_{n\to\infty} d(b_n, c_n).$$

Hence,

$$d^*([\mathbf{a}], [\mathbf{c}]) \leq d^*([\mathbf{a}], [\mathbf{b}]) + d^*([\mathbf{b}], [\mathbf{c}]),$$

as required.

Since (M1)–(M3) are satisfied, d^* is a metric on X^*.

Chapter 24

Fractals

Introduction

In this, the final chapter of Book F and of the module, we show how the theory of metric spaces is used to great effect in the study of fractals, an area where there is much current interest and research activity. You have probably seen computer-generated pictures of fractals and may even have produced such pictures yourself. One picture you may have seen is the *von Koch snowflake* (Figure 0.1).

There is no standard definition of a fractal, but there are certain characteristics that we expect a fractal to have. The von Koch snowflake is typical of many fractals in the following ways:

- it has a fine structure – the closer you look at it, the more detail you see
- it has some degree of self-similarity – the same patterns keep appearing
- it is difficult to describe in terms of classical geometry.

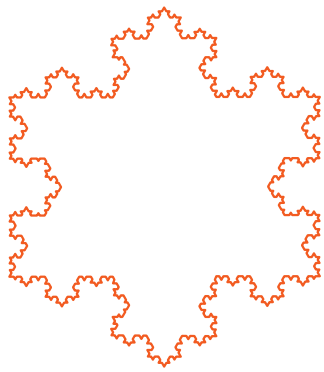

Figure 0.1

Despite having an intricate structure, many fractals can be defined very simply. In this chapter we use the Contraction Mapping Theorem of Chapter 23 to generate a large class of fractals, including the von Koch snowflake.

We also look at the *dimensions* of fractals. Many fractals cannot be described adequately in terms of the usual integer-valued definitions of dimension. For example, the von Koch snowflake does not have finite length (suggesting that it has dimension greater than 1) and has zero area (suggesting that it has dimension less than 2). We introduce a definition of dimension that gives a good indication of the 'size' of many fractals, but which can take non-integer values.

Chapter guide

In Section 1 we look at three examples of fractals and derive some of their properties. In Section 2, we look at how we may define these fractals as fixed points of contraction mappings on the space of non-empty compact subsets of either the real line or the plane. This requires introducing a way of measuring the distance between compact sets, the Hausdorff metric.

In the last section, we give a brief introduction to the basic ideas behind fractal dimension and show that for certain fractal sets there is an easy way to calculate their dimension.

1 Examples of fractals

In this section we look at some well-known examples of fractals. By studying their properties, you should get a feeling for the characteristics that are typical of fractals in general.

1.1 Cantor sets

We begin by looking at a class of fractals known as *Cantor sets*. These are subsets of the real line that are obtained by systematically removing open intervals from a closed interval, to obtain a compact, totally disconnected set. The best known of these is the *middle-third Cantor set* that you met earlier (see Figure 1.1).

Figure 1.1 Middle-third Cantor set

You may recall that we obtain the middle-third Cantor set by a simple recursive procedure. We begin with the line segment $K_0 = [0, 1]$ (Figure 1.2).

Figure 1.2

This explains the name *middle-third* Cantor set.

We then remove the open middle third of K_0 – the open interval $(\frac{1}{3}, \frac{2}{3})$ – leaving a set K_1 made up of two closed intervals, each of length one-third. So

$$K_1 = \left[0, \tfrac{1}{3}\right] \cup \left[\tfrac{2}{3}, 1\right],$$

as in Figure 1.3.

Figure 1.3

We then apply the same procedure to each of the two line segments in K_1 to give the set

$$K_2 = \left[0, \tfrac{1}{9}\right] \cup \left[\tfrac{2}{9}, \tfrac{1}{3}\right] \cup \left[\tfrac{2}{3}, \tfrac{7}{9}\right] \cup \left[\tfrac{8}{9}, 1\right],$$

as in Figure 1.4.

Figure 1.4

We continue in this way, constructing a set K_n for each $n \in \mathbb{N}$. Thus K_n is formed by removing the open middle third of each line segment in K_{n-1}. The set K_n consists of 2^n closed intervals, each with length $\frac{1}{3^n}$.

The middle-third Cantor set is the intersection of the sets K_n. We denote this set by

$$K_{\mathrm{C}} = \bigcap_{n \in \mathbb{N}} K_n,$$

as in Figure 1.5.

You will see that many fractals can be defined by a similar recursive method.

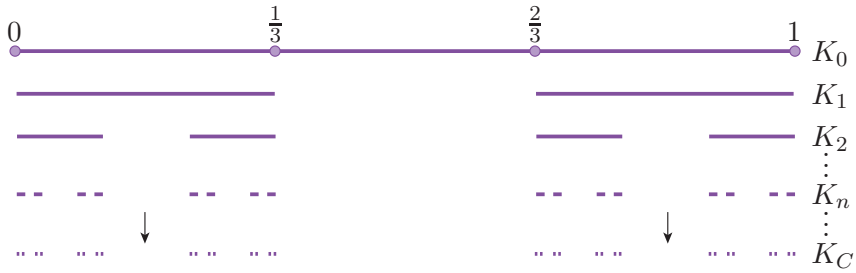

Figure 1.5

You may wonder if K_{C} contains *any* points at all! However, each set K_n is closed and bounded, and hence compact. Since $K_{n+1} \subseteq K_n$ for each $n \in \mathbb{N}$, the Cantor set K_{C} is the intersection of a nested sequence of non-empty compact sets and is therefore *non-empty* and compact. In fact, K_{C} contains the endpoints of each of the 2^n line segments forming K_n, for each $n \in \mathbb{N}$, since none of these endpoints is removed in the recursive procedure. Thus K_{C} is an *infinite* set. In fact, K_{C} contains *uncountably* many points, as we now show.

See Chapter 22, Theorem 6.3.

Theorem 1.1

The middle-third Cantor set K_{C} is uncountable.

Proof We suppose that K_{C} is countable and obtain a contradiction.

If K_{C} is countable, then we can list all its elements, say

$$K_{\mathrm{C}} = \{c_1, c_2, c_3, \ldots\}. \tag{1}$$

Now c_1 is not in both of the intervals $\left[0, \frac{1}{3}\right]$, $\left[\frac{2}{3}, 1\right]$, so we may let I_1 be whichever it is of these two intervals that does not contain c_1.

Let I_1' be I_1 with its middle third removed, so I_1' consists of two disjoint, closed intervals, each of length $\frac{1}{9}$. As c_2 is not in both of these intervals, we may let I_2 be whichever it is of these two intervals that does not contain c_2.

Next, let I_2' be I_2 with its middle third removed, so I_2' consists of two disjoint, closed intervals, each of length $\frac{1}{27}$. As c_3 is not in both of these intervals, we may let I_3 be whichever it is of these two intervals that does not contain c_3.

Continuing in this way, we obtain

$$K_0 \supseteq I_1 \supseteq I_2 \supseteq I_3 \supseteq \cdots \supseteq I_k \supseteq \cdots$$

where each I_k is a closed interval of length $\frac{1}{3^k}$. In particular, each I_k is closed and bounded, so is compact. Thus

$$\bigcap_{j=1}^{\infty} I_j$$

is the intersection of a nested sequence of non-empty, compact sets, and is therefore non-empty and compact. So there is a point $c \in \cap_{j=1}^{\infty} I_j \subseteq \bigcap_{j=1}^{\infty} K_j = K_{\mathrm{C}}$, yet $c \neq c_k$ for any k, since $c_k \notin I_k$, and hence $c_k \notin \cap_{j=1}^{\infty} I_j$ for all k.

See Chapter 22, Theorem 6.3.

Thus we have found a point $c \in K_{\mathrm{C}}$ with $c \notin \{c_1, c_2, c_3, \ldots\}$. This contradicts (1), indicating that our original supposition was false, and K_{C} is uncountable. ■

Our study has revealed that K_{C} is an infinite set with uncountably many points. It is often the case with fractals that natural questions lead to interesting, perhaps even unexpected, results. One such question is: '*What is the dimension of* K_{C}*?*'

Exercise 1.1

For each $n \in \mathbb{N}$, calculate the total length L_n of the line segments that form K_n, and find $\lim_{n\to\infty} L_n$.

Since K_{C} is a subset of K_n for each $n \in \mathbb{N}$, the solution to Exercise 1.1 suggests that K_{C} has 'zero length', and is therefore 'smaller' than a one-dimensional set (such as a line segment). It is also 'larger' than a zero-dimensional set (such as a point), since it contains infinitely many points. In Section 3, we give a definition of dimension that allows us to assign K_{C} a dimension between 0 and 1.

Self-similarity

In the construction of the middle-third Cantor set, the set K_1 is made up of two copies of K_0, each scaled by $\frac{1}{3}$ (Figure 1.6).

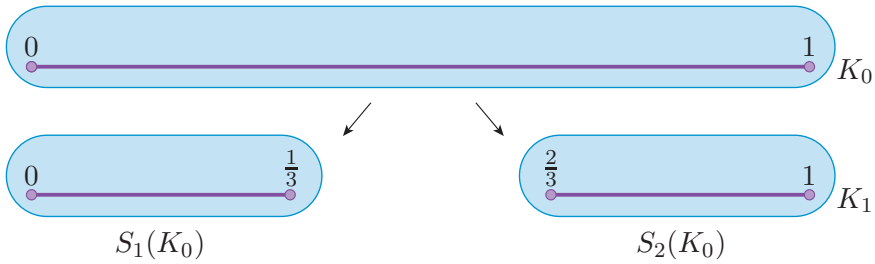

Figure 1.6

Similarly, the set K_2 consists of two copies of K_1, each scaled by $\frac{1}{3}$ – and in general, for each $n \in \mathbb{N}$, the set K_{n+1} consists of two copies of K_n, each scaled by $\frac{1}{3}$.

The mappings taking K_n to these two smaller copies of itself are the functions $S_1, S_2 \colon \mathbb{R} \to \mathbb{R}$, defined by

$$S_1(x) = \tfrac{1}{3}x \quad \text{and} \quad S_2(x) = \tfrac{1}{3}(x+2).$$

The mapping S_1 scales the Cantor set by $\frac{1}{3}$, and the mapping S_2 scales the Cantor set by $\frac{1}{3}$ and translates the resulting set $\frac{2}{3}$ to the right.

Each of the mappings S_1 and S_2 is a *similarity*.

Definition 1.2 *Similarity*

A mapping $S \colon \mathbb{R}^m \to \mathbb{R}^m$ is a **similarity** if there is a $c > 0$ such that

$$d^{(m)}(S(\mathbf{x}), S(\mathbf{y})) = c\, d^{(m)}(\mathbf{x}, \mathbf{y}) \quad \text{for all } \mathbf{x}, \mathbf{y} \in \mathbb{R}^m.$$

The constant c is the **similarity ratio**.

Each of the functions S_1 and S_2 scales by $\frac{1}{3}$, so has similarity ratio $\frac{1}{3}$. Any translation, reflection or rotation, or more generally any isometry, is a similarity with similarity ratio 1, because distances remain unchanged under these transformations.

Isometries were discussed in Chapter 23.

Exercise 1.2

Write down the similarity ratio of each of the following mappings.

(a) $S \colon \mathbb{R} \to \mathbb{R}$ defined by $S(x) = \frac{1}{2}x + 3$

(b) $S \colon \mathbb{R}^2 \to \mathbb{R}^2$ defined by $S((x, y)) = (\frac{1}{4}y, \frac{1}{4}(x+3))$

For the similarities S_1 and S_2 used to define the Cantor set, each set K_n in the construction of K_C satisfies

$$K_{n+1} = S_1(K_n) \cup S_2(K_n).$$

Since K_C is a 'limit' of the sets K_n, we may expect that

$$K_C = S_1(K_C) \cup S_2(K_C),$$

as in Figure 1.7.

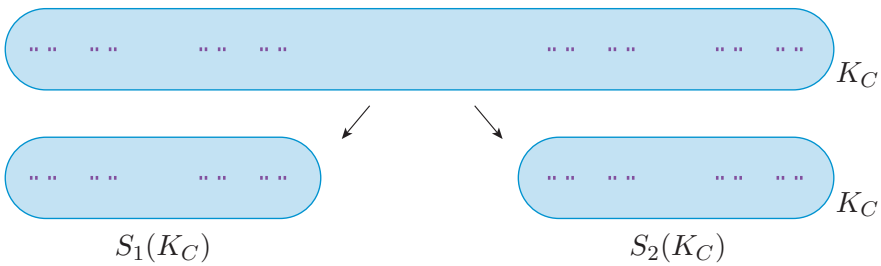

Figure 1.7

In Section 2, we show that this is indeed the case. We say that the middle-third Cantor set is *self-similar*.

The set A consists of k copies of itself, each of which is the image of A under a similarity.

Definition 1.3 *Self-similar sets*

A set $A \subseteq \mathbb{R}^m$ is **self-similar** if there are similarities $S_i \colon \mathbb{R}^m \to \mathbb{R}^m$ with similarity ratios $0 < c_i < 1$, for $1 \leq i \leq k$, such that

$$A = \bigcup_{i=1}^{k} S_i(A).$$

The self-similarity of the middle-third Cantor set means that it is regular, with the same patterns (but smaller) appearing again and again the closer you look at the set.

Most fractals have some degree of self-similarity. There are also many simple sets that are self-similar.

Exercise 1.3

Show that the line segment $[0, 1]$ is self-similar, by writing down two similarities $S_1, S_2 \colon \mathbb{R} \to \mathbb{R}$, with similarity ratios $0 < c_1 < 1$ and $0 < c_2 < 1$, such that

$$[0, 1] = S_1([0, 1]) \cup S_2([0, 1]).$$

Not all simple sets are self-similar; for example, it can be shown that no circle is self-similar.

Many other fractals can be obtained by a procedure similar to that used to obtain the middle-third Cantor set. We now ask you to investigate one of these.

Exercise 1.4

Let $K_0 = [0, 1]$. Remove the second and fourth quarters of K_0 to give a new set

$$K_1 = [0, \tfrac{1}{4}] \cup [\tfrac{1}{2}, \tfrac{3}{4}],$$

as in Figure 1.8.

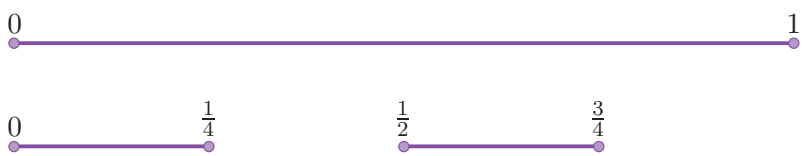

Figure 1.8

For each $n \in \mathbb{N}$, apply this procedure to each closed line segment in K_n to give a new set K_{n+1}.

(a) Draw the set K_2.

(b) For each $n \in \mathbb{N}$, calculate the total length L_n of the line segments that form K_n, and find $\lim_{n\to\infty} L_n$.

(c) Write down similarities $S_1, S_2 \colon \mathbb{R} \to \mathbb{R}$ such that

$$K_{n+1} = S_1(K_n) \cup S_2(K_n) \quad \text{for each } n \in \mathbb{N},$$

and state their similarity ratios.

1.2 The Sierpiński gasket

Our next example of a fractal is the *Sierpiński gasket* (Figure 1.9).

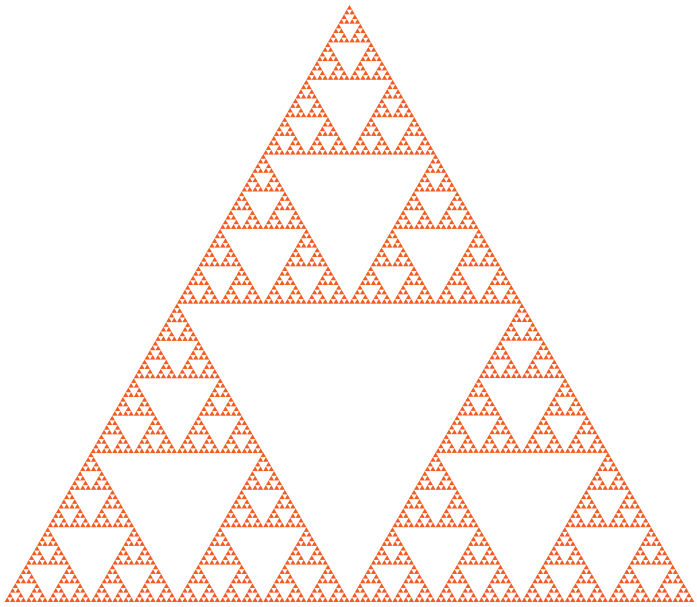

Figure 1.9

Wacław Sierpiński (1882–1969) was one of the founders of the Polish school of mathematics, which focused on point-set topology (a subject involving the study of continuous functions in an even more general context than metric spaces), real analysis and logic. He published more than 600 papers, and introduced a number of thought-provoking sets, including the 'gasket' that is named after him.

This fractal is obtained by a recursive procedure. We begin (Figure 1.10) with an upright closed equilateral triangle K_0 whose base is the line segment joining the points $(0, 0)$ and $(1, 0)$.

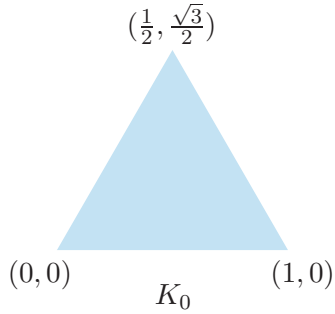

Figure 1.10

We then remove an open equilateral triangle from K_0, leaving a set K_1 made up of three closed equilateral triangles, each with height half that of K_0 (Figure 1.11).

We apply this procedure to each of the three triangles in K_1 to give a new set K_2 (Figure 1.12).

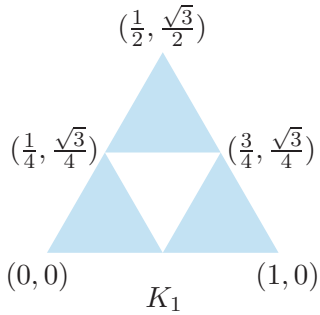

Figure 1.11

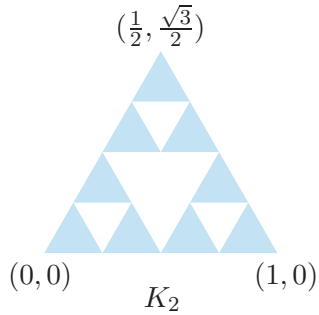

Figure 1.12

We continue in this way, constructing a set K_n for each $n \in \mathbb{N}$. The set K_n is formed from 3^n equilateral triangles, each with height $1/2^n$ times the height of K_0. The Sierpiński gasket K_{G} is defined to be the intersection of the sets K_n,

$$K_{\mathrm{G}} = \bigcap_{n \in \mathbb{N}} K_n.$$

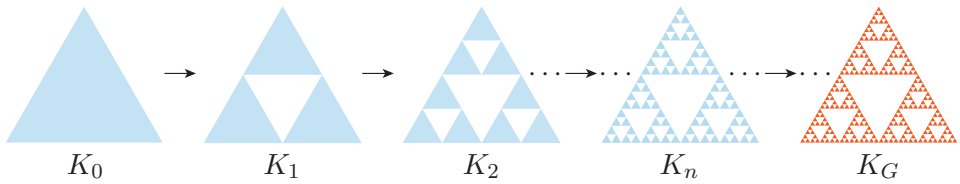

Figure 1.13

Each set K_n is closed and bounded, and hence compact. Since $K_{n+1} \subseteq K_n$, for each $n \in \mathbb{N}$, the Sierpiński gasket K_{G} is the intersection of a nested sequence of non-empty compact sets, and is therefore non-empty and compact. In fact, K_{G} contains the boundary of each of the 3^n equilateral triangles forming K_n, for each $n \in \mathbb{N}$, since no boundary points are removed in the recursive procedure (Figure 1.13).

See Chapter 22, Theorem 6.3.

Exercise 1.5

For each $n \in \mathbb{N}$, calculate the total length of the boundaries of the 3^n equilateral triangles that form K_n. Deduce that, given any $N \in \mathbb{R}$, there is a finite collection of equilateral triangles with the property that the Sierpiński gasket contains their boundaries, *and* the total length of their boundaries exceeds N.

We now ask you to show that K_{G} has 'zero area'.

Exercise 1.6

For each $n \in \mathbb{N}$, calculate the area of K_n. Deduce that the Sierpiński gasket has zero area.

The Sierpiński gasket seems to be smaller than a two-dimensional set (such as a rectangle), since it has zero area, and yet it seems to be larger than a one-dimensional set (such as a line segment), since it is bounded but does not have finite length.

We address this issue in Section 3.

Note that the Sierpiński gasket appears to be self-similar – it appears to be made up of three copies of itself, each with height $\frac{1}{2}$ times the height of the whole gasket.

Exercise 1.7

Let K_0 and K_1 be the first two sets in the construction of the Sierpiński gasket. Write down three similarities $S_1, S_2, S_3 \colon \mathbb{R}^2 \to \mathbb{R}^2$ such that

$$K_1 = S_1(K_0) \cup S_2(K_0) \cup S_3(K_0),$$

and state their similarity ratios.

For three such similarities S_1, S_2 and S_3,

$$K_{n+1} = S_1(K_n) \cup S_2(K_n) \cup S_3(K_n),$$

for each $n \in \mathbb{N}$, and this suggests that

$$K_{\mathrm{G}} = S_1(K_{\mathrm{G}}) \cup S_2(K_{\mathrm{G}}) \cup S_3(K_{\mathrm{G}}).$$

We justify this in Section 2.

Many other fractals can be obtained by a procedure similar to that used to obtain the Sierpiński gasket.

Exercise 1.8

Let K_0 be the closed square with corners at the points $(0,0)$, $(0,1)$, $(1,0)$ and $(1,1)$. Divide K_0 into nine subsquares and remove all but the middle and corner closed subsquares to give a new set K_1 (Figure 1.14).

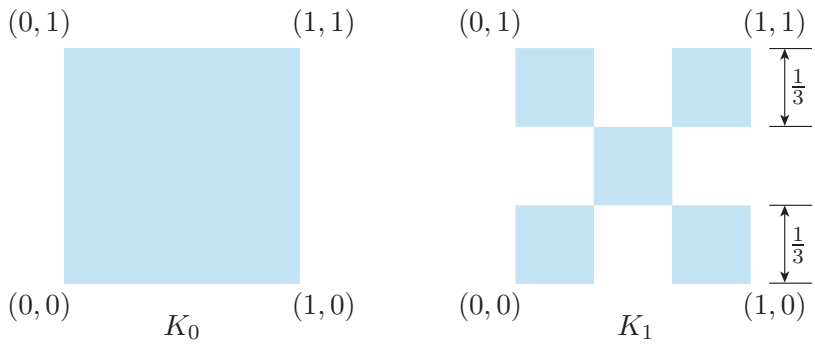

Figure 1.14

For each $n \in \mathbb{N}$, apply this procedure to each of the closed squares in K_n to give a new set K_{n+1}. Let K denote the intersection of the sets K_n.

(a) Draw the set K_2.

(b) Show that K is non-empty and compact and has zero area.

(c) Write down five similarities $S_1, S_2, \ldots, S_5 \colon \mathbb{R}^2 \to \mathbb{R}^2$ such that

$$K_{n+1} = \bigcup_{i=1}^{5} S_i(K_n) \quad \text{for each } n \in \mathbb{N},$$

and state their similarity ratios.

1.3 The von Koch curve

We conclude this section by returning to the von Koch snowflake that you met in the Introduction. The snowflake is made up of three copies of the von Koch curve, as shown in Figure 1.15.

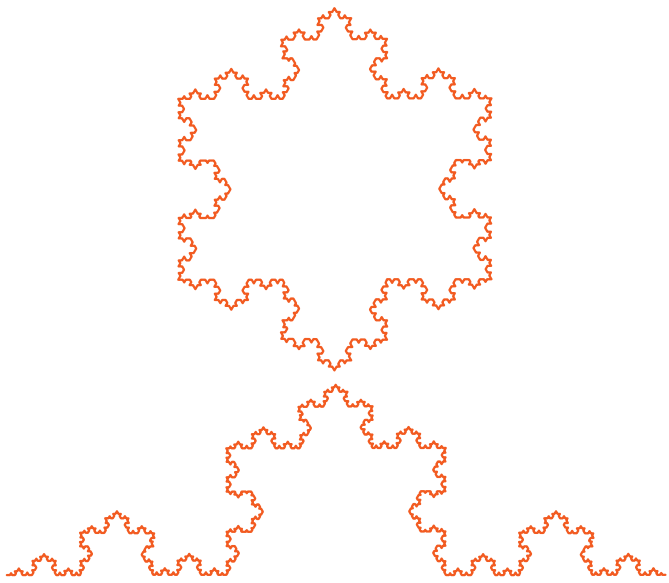

Figure 1.15

The Swedish mathematician Helge von Koch (1870–1924) described this curve in 1906, showing that no point on the curve has a tangent.

Although the von Koch curve appears to be very intricate, it is obtained by a simple procedure. We begin with the line segment $K_0 = [0,1] \times \{0\}$ (Figure 1.16).

Figure 1.16

We then remove the middle third of K_0 and replace it with the other two sides of the equilateral triangle based on the missing segment. This gives a set K_1 (Figure 1.17).

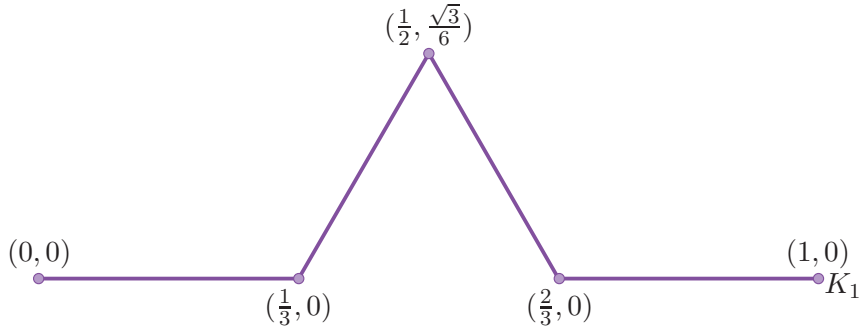

Figure 1.17

We then apply this procedure to each of the four line segments of K_1 to give a new set K_2 (Figure 1.18).

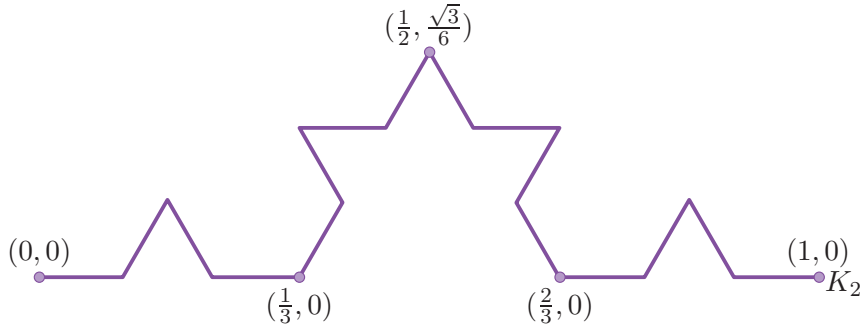

Figure 1.18

We continue in this way, constructing sets K_n for each $n \in \mathbb{N}$. We wish to define the von Koch curve K_{V} to be the limit of the sets K_n as $n \to \infty$ (Figure 1.19 overleaf).

At this point it is not clear how to define the *limit* of the sets K_n as $n \to \infty$. In each of our previous examples, we had a nested sequence of compact sets, and so could define the limit set to be the intersection of these nested sets – this must be a non-empty compact set. In this case, the sets K_n are compact, but they are not nested, and it is not clear that a well-defined limit exists.

In Section 2 we show that this limit does exist.

Although it is not obvious how to define the limit of the sets K_n, we can identify certain points that must belong to the von Koch curve. For each $n \in \mathbb{N}$, the endpoints of each line segment in K_n are not removed by the subsequent stages of the construction, and so each of these endpoints belongs to K_{V}.

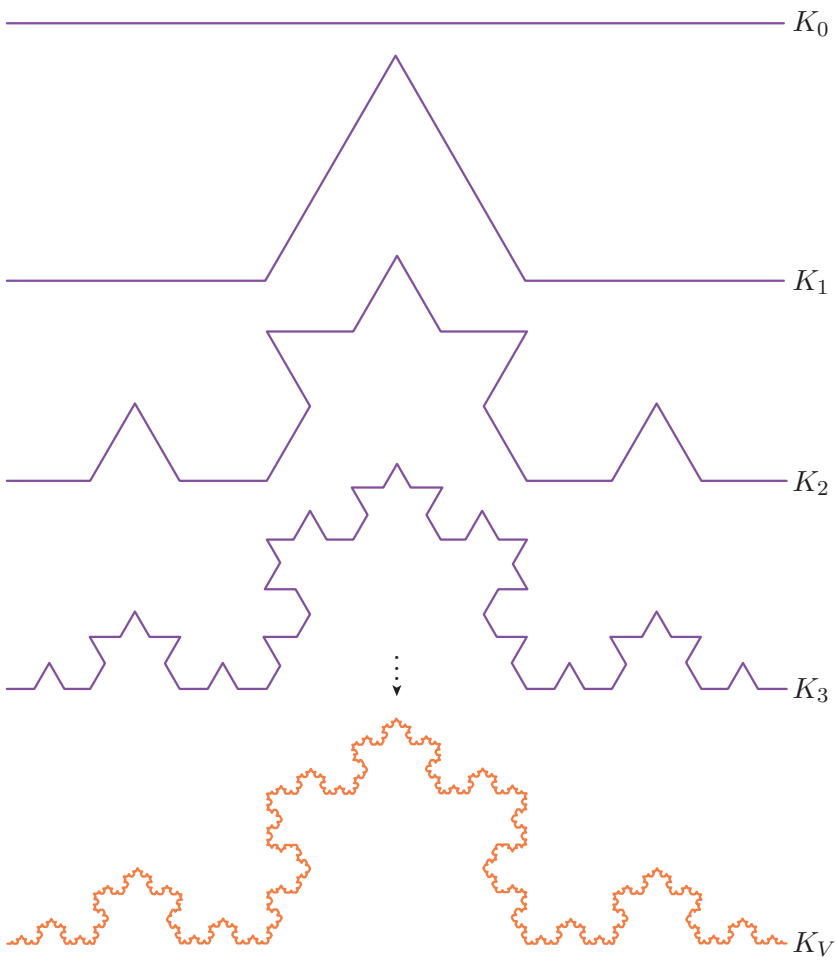

Figure 1.19

Exercise 1.9

For each $n \in \mathbb{N}$, determine the length of each line segment that makes up K_n.

Using the fact that the length of the section of the von Koch curve lying between the two endpoints of one of these line segments cannot be less than the distance between the two endpoints, deduce that K_V does not have finite length.

We have just seen that the von Koch curve is very different from most familiar curves. It lies within a bounded part of the plane, yet it is so 'wiggly' that it does not have finite length. The 'wiggly' nature of the curve also means that no point on the curve has a tangent. It is thus difficult to describe the curve in terms of classical geometric shapes.

This is typical of fractals.

For each $n \in \mathbb{N}$, the set K_{n+1} is made up of four copies of K_n, each scaled by $\frac{1}{3}$. So

$$K_{n+1} = S_1(K_n) \cup S_2(K_n) \cup S_3(K_n) \cup S_4(K_n), \tag{2}$$

where the mappings S_1, S_2, S_3, S_4 are described below.

In fact, each of the mappings S_1, S_2, S_3, S_4 is a similarity: you should be able to write down similarity ratios for each of them after reading their descriptions.

S_1 scales by a factor of $\frac{1}{3}$, as in Figure 1.20.

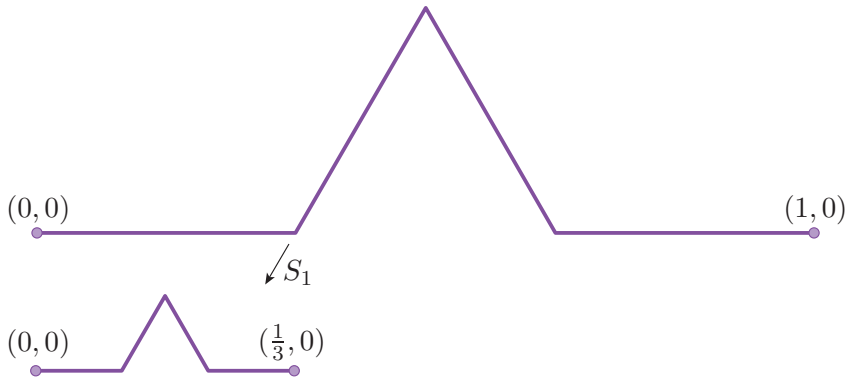

Figure 1.20

S_2 scales by a factor of $\frac{1}{3}$, rotates anticlockwise by $\frac{\pi}{3}$ about the point $(0, 0)$, and translates by $(\frac{1}{3}, 0)$, as in Figure 1.21.

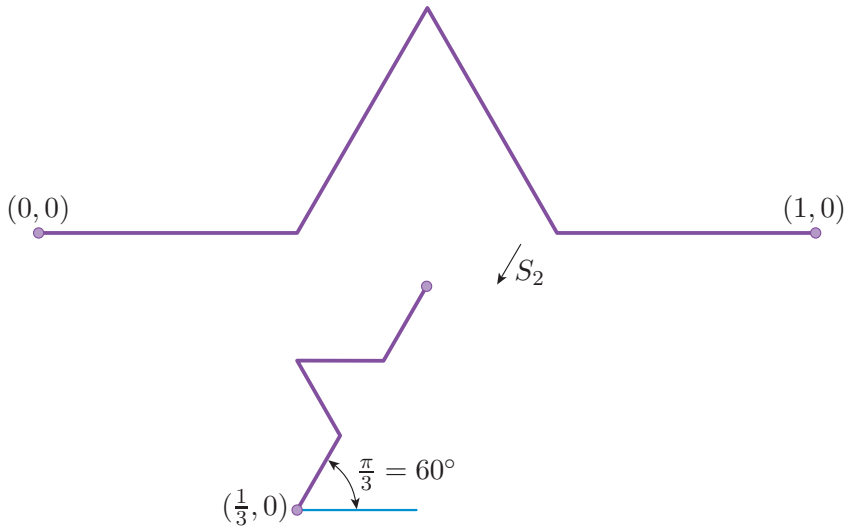

Figure 1.21

S_3 scales by a factor of $\frac{1}{3}$, rotates clockwise by $\frac{\pi}{3}$ about the point $(0, 0)$, and translates by $(\frac{1}{2}, \frac{\sqrt{3}}{6})$, as in Figure 1.22.

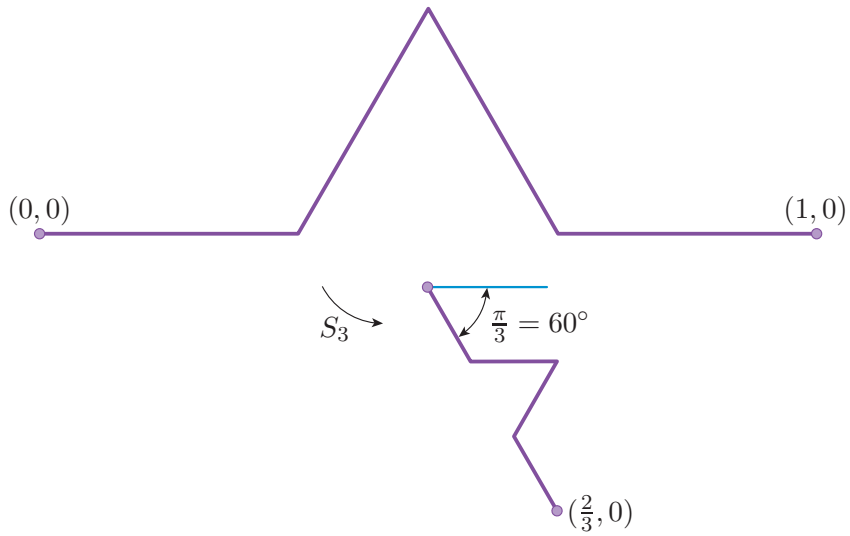

Figure 1.22

S_4 scales by a factor of $\frac{1}{3}$ and translates by $(\frac{2}{3}, 0)$, as in Figure 1.23.

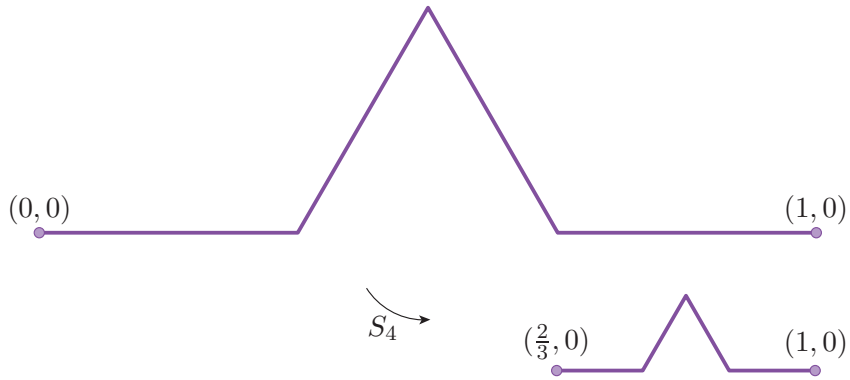

Figure 1.23

Equation (2) suggests that if K_{V} is the limit of the sets K_n, then

This is justified in Section 2.

$$K_{\mathrm{V}} = S_1(K_{\mathrm{V}}) \cup S_2(K_{\mathrm{V}}) \cup S_3(K_{\mathrm{V}}) \cup S_4(K_{\mathrm{V}}).$$

So the von Koch curve appears to be self-similar.

You have now met some well-known fractals: the middle-third Cantor set, the Sierpiński gasket and the von Koch curve. We will refer to these examples throughout the chapter.

2 Using the Contraction Mapping Theorem

In the previous section, you met several examples of fractals, each of which appears to be self-similar. In other words, for each such fractal K, there appears to be a finite family of similarities, $S_1, S_2, \ldots, S_k$, such that

$$K = S_1(K) \cup S_2(K) \cup \cdots \cup S_k(K) = \bigcup_{i=1}^{k} S_i(K).$$

Each fractal K is constructed by taking a compact starting set K_0 and repeatedly applying the similarities $S_1, S_2, \ldots, S_k$.

In this section, we use the Contraction Mapping Theorem to show that if we take *any* non-empty compact set K_0 and *any* family of contraction mappings $S_1, S_2, \ldots, S_k$, then this recursive procedure defines a non-empty compact set K satisfying $K = \bigcup_{i=1}^{k} S_i(K)$. By looking at different families of contraction mappings, we can construct many different fractals. Note that all the similarities in Section 1 are in fact contraction mappings.

2.1 The Hausdorff metric

We begin by reminding you of the definition of a contraction mapping, and the Contraction Mapping Theorem.

Definition 2.1 *Definition 3.3 of Chapter 23*

Let (X, d) be a metric space. A function $T\colon X \to X$ is a **contraction mapping** if there is a real number $0 \leq \lambda < 1$ such that

$$d(T(x), T(y)) \leq \lambda d(x, y) \quad \text{for all } x, y \in X.$$

Any such real number λ is a **contraction ratio** for T.

Notice that if λ is a contraction ratio for T and $\lambda < \mu < 1$ then μ is also a contraction ratio for T.

Theorem 2.2 *Theorem 3.6 of Chapter 23 (Contraction Mapping Theorem)*

Let (X, d) be a complete metric space, and let $T\colon X \to X$ be a contraction mapping. Then there is a unique element $x_T \in X$ such that

$$T(x_T) = x_T.$$

Moreover, for each $x \in X$, the sequence of iterates

$$x,\ T(x),\ T^2(x),\ T^3(x),\ \ldots$$

converges to this unique fixed point x_T – that is, for any $x \in X$,

$$d(T^n(x), x_T) \to 0 \text{ as } n \to \infty.$$

We shall use this result to generate a large class of fractals, each of which is the fixed point x_T of a suitable contraction mapping T. We begin by defining a suitable metric space (X, d).

Recall that the compact subsets of $\mathbb{R}^m$ are those that are closed and bounded.

In each of our earlier examples, we repeatedly applied mappings to a *compact* subset of $\mathbb{R}^m$ (where m was 1 or 2). Since this process generated fractals that are also compact subsets of $\mathbb{R}^m$, it seems reasonable to take our set X to be a set whose members are compact subsets of $\mathbb{R}^m$.

Definition 2.3

For each $m \in \mathbb{N}$, the set $\mathcal{K}(\mathbb{R}^m)$ is defined as follows:

$$\mathcal{K}(\mathbb{R}^m) = \{K \subseteq \mathbb{R}^m : K \text{ is non-empty and compact}\}.$$

Each member of $\mathcal{K}(\mathbb{R}^m)$ is *itself* a set. For example, the *set* $[0, 1] \subseteq \mathbb{R}$ is non-empty and compact, so it is a member of $\mathcal{K}(\mathbb{R})$. Thus we write $[0, 1] \in \mathcal{K}(\mathbb{R})$.

We are usually interested in the sets $\mathcal{K}(\mathbb{R})$ and $\mathcal{K}(\mathbb{R}^2)$. For example, the middle-third Cantor set is a set in $\mathcal{K}(\mathbb{R})$ that is generated by repeatedly applying mappings to sets in $\mathcal{K}(\mathbb{R})$, and the Sierpiński gasket is a set in $\mathcal{K}(\mathbb{R}^2)$ that is generated by repeatedly applying mappings to sets in $\mathcal{K}(\mathbb{R}^2)$.

Exercise 2.1

Which of the following sets belong to $\mathcal{K}(\mathbb{R})$?

(a) $(-1, 1]$ (b) $[0, 2]$ (c) $[0, \infty)$ (d) $\{1\}$ (e) $[0, 2] \cap [4, 6]$

In order to apply the Contraction Mapping Theorem, we need to define a *metric* on the set $\mathcal{K}(\mathbb{R}^m)$ – that is, a distance between any two sets in $\mathcal{K}(\mathbb{R}^m)$ for which the metric properties (M1)–(M3) are satisfied. We begin by considering the simpler problem of how to define the distance between a *point* in $\mathbb{R}^m$ and a set K in $\mathcal{K}(\mathbb{R}^m)$. It seems reasonable to define this as follows.

Definition 2.4 *Euclidean distance between a point and a set*

Let $\mathbf{x} \in \mathbb{R}^m$ and let $K \in \mathcal{K}(\mathbb{R}^m)$, for some $m \in \mathbb{N}$.

The **Euclidean distance** between the point $\mathbf{x}$ and the set K is

$$d^{(m)}(\mathbf{x}, K) = \min\{d^{(m)}(\mathbf{x}, \mathbf{k}) : \mathbf{k} \in K\}.$$

Remark

This definition makes sense since the function $d^{(m)}(\mathbf{x}, \mathbf{k})$ is a *continuous* function of $\mathbf{k}$ on the *compact* set K. It follows from the General Extreme Value Theorem that there is a $\mathbf{c} \in K$ such that

$$d^{(m)}(\mathbf{x}, \mathbf{c}) \leq d^{(m)}(\mathbf{x}, \mathbf{k}) \quad \text{for each } \mathbf{k} \in K.$$

So

$$d^{(m)}(\mathbf{x}, K) = d^{(m)}(\mathbf{x}, \mathbf{c}).$$

See Book D, Chapter 14, Exercise 1.10 for a proof of this result when $\mathbf{x} = \mathbf{0}$.

See Chapter 22, Theorem 2.13.

For example, consider the point $\mathbf{x} = (3, 0)$ and the set $K = \{(x, y) : |x| \leq 2 \text{ and } |y| \leq 1\}$, which is a member of $\mathcal{K}(\mathbb{R}^2)$ (Figure 2.1).

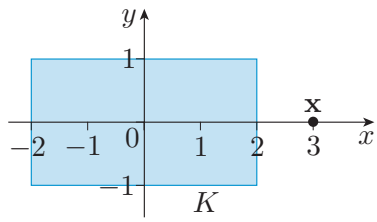

Figure 2.1

The closest point in K to the point $(3, 0)$ is the point $(2, 0)$, and so

$$d^{(2)}((3, 0), K) = d^{(2)}((3, 0), (2, 0)) = 1.$$

Exercise 2.2

Determine the following.

(a) $d^{(1)}(1, [2, 3])$ (b) $d^{(1)}(1, [0, 2])$

(c) $d^{(2)}((0, 3), \{(x, y) : x^2 + y^2 \leq 1\})$ (d) $d^{(2)}((2, 1), \{(3, y) : |y| \leq 1\})$

Note that $d^{(m)}(\mathbf{x}, K) = 0$ precisely when there is a point $\mathbf{c} \in K$ such that $d^{(m)}(\mathbf{x}, \mathbf{c}) = 0$ – that is, $\mathbf{x} = \mathbf{c}$. This gives the following result.

Lemma 2.5

Let $\mathbf{x} \in \mathbb{R}^m$ and let $K \in \mathcal{K}(\mathbb{R}^m)$, for some $m \in \mathbb{N}$. Then

$d^{(m)}(\mathbf{x}, K) = 0$ if, and only if, $\mathbf{x} \in K$.

The following lemma gives a useful property of $d^{(m)}$ for compact sets.

Lemma 2.6

Let $\mathbf{x} \in \mathbb{R}^m$ and let $J, K \in \mathcal{K}(\mathbb{R}^m)$, for some $m \in \mathbb{N}$. Let $J \subseteq K$; then

$$d^{(m)}(\mathbf{x}, K) \leq d^{(m)}(\mathbf{x}, J).$$

Exercise 2.3

Prove Lemma 2.6.

Having established this definition of the distance from a point to a compact set, we now extend it to the definition of the distance from one compact set to another as follows.

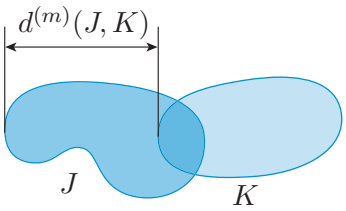

Figure 2.2

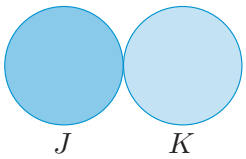

Figure 2.3

Definition 2.7 *Euclidean distance between compact sets*

Let $J, K \in \mathcal{K}(\mathbb{R}^m)$, for some $m \in \mathbb{N}$. The **Euclidean distance** between the set J and the set K (see Figure 2.2) is

$$d^{(m)}(J,K) = \max\{d^{(m)}(\mathbf{j},K) : \mathbf{j} \in J\}.$$

Remark

We use 'max' here instead of 'min' as 'min' would give a distance of zero between the pairs of sets shown in Figures 2.2 and 2.3.

We want distinct sets to be a positive distance apart, otherwise we do not have a metric.

See Chapter 22, Theorem 2.13.

This definition makes sense since the function $d^{(m)}(\mathbf{j},K)$ is a *continuous* function of $\mathbf{j}$ on the *compact* set J. It follows from the General Extreme Value Theorem that there is a $\mathbf{d} \in J$ such that

$$d^{(m)}(\mathbf{j},K) \leq d^{(m)}(\mathbf{d},K) \quad \text{for each } \mathbf{j} \in J.$$

So

$$d^{(m)}(J,K) = d^{(m)}(\mathbf{d},K).$$

Since there is a $\mathbf{c} \in K$ such that $d^{(m)}(\mathbf{d},K) = d^{(m)}(\mathbf{d},\mathbf{c})$, it follows that

$$d^{(m)}(J,K) = d^{(m)}(\mathbf{d},\mathbf{c}).$$

In other words, the Euclidean distance between two non-empty compact sets J and K is realised by a pair of points, one from J and one from K.

In the exercises, we ask you to prove that $d^{(m)}(\mathbf{j},K)$ is a continuous function of $\mathbf{j}$.

Worked Exercise 2.8

Let

$$J = \{(3,y) : |y| \leq 1\} \quad \text{and} \quad K = \{(x,y) : |x| \leq 2 \text{ and } |y| \leq 1\}.$$

Determine $d^{(2)}(J,K)$ and $d^{(2)}(K,J)$.

Solution

We begin by sketching the sets J and K.

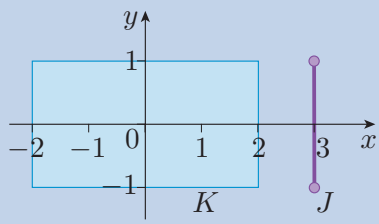

Figure 2.4

The points in J are all at distance 1 from the set K; for example,

$$d^{(2)}(J,K) = d^{(2)}((3,0),K) = d^{(2)}((3,0),(2,0)) = 1.$$

The points in K that are furthest from the set J are those on the left-hand edge of K. One such point is $(-2,0)$, and so

$$d^{(2)}(K,J) = d^{(2)}((-2,0),J) = d^{(2)}((-2,0),(3,0)) = 5.$$

Exercise 2.4

Let

$$J = \{(x,y) : x^2 + y^2 \leq 1\} \quad \text{and} \quad K = \{(1,0)\}.$$

Determine $d^{(2)}(J,K)$ and $d^{(2)}(K,J)$.

These examples show that the Euclidean distance from one non-empty compact set to another is *not* a metric, since the symmetry property (M2) is not satisfied. In fact, (M1) is not satisfied either, as the first part of the following exercise shows.

You may suspect that we have chosen a poor definition of distance between compact sets, but you will see shortly how to rescue this definition.

Exercise 2.5

Let $J, K \in \mathcal{K}(\mathbb{R}^m)$, for some $m \in \mathbb{N}$, and let $J \subseteq K$.

(a) Show that $d^{(m)}(J,K) = 0$.

(b) Show that $d^{(m)}(L,J) \geq d^{(m)}(L,K)$, for each $L \in \mathcal{K}(\mathbb{R}^m)$.

We now show that the Euclidean distance *does* satisfy the Triangle Inequality – property (M3). This suggests that, by modifying the definition of this distance slightly, we may be able to obtain a definition of distance between two compact sets that is a metric, as we require.

Lemma 2.9

Let $J, K, L \in \mathcal{K}(\mathbb{R}^m)$, for some $m \in \mathbb{N}$. Then

$$d^{(m)}(J,L) \leq d^{(m)}(J,K) + d^{(m)}(K,L).$$

Proof We know that $d^{(m)}$ is a metric on $\mathbb{R}^m$.

By the Triangle Inequality for $d^{(m)}$, for each $\mathbf{j} \in J$, $\mathbf{k} \in K$, $\mathbf{l} \in L$,

$$d^{(m)}(\mathbf{j},\mathbf{l}) \leq d^{(m)}(\mathbf{j},\mathbf{k}) + d^{(m)}(\mathbf{k},\mathbf{l}).$$

Taking $\mathbf{l} \in L$ such that $d^{(m)}(\mathbf{k},\mathbf{l}) = d^{(m)}(\mathbf{k},L)$, we obtain

$$d^{(m)}(\mathbf{j},\mathbf{l}) \leq d^{(m)}(\mathbf{j},\mathbf{k}) + d^{(m)}(\mathbf{k},L).$$

Taking $\mathbf{k} \in K$ such that $d^{(m)}(\mathbf{j}, \mathbf{k}) = d^{(m)}(\mathbf{j}, K)$, we obtain

$$d^{(m)}(\mathbf{j}, \mathbf{l}) \leq d^{(m)}(\mathbf{j}, K) + d^{(m)}(\mathbf{k}, L).$$

Since $d^{(m)}(\mathbf{j}, L) \leq d^{(m)}(\mathbf{j}, \mathbf{l})$ and $d^{(m)}(\mathbf{k}, L) \leq d^{(m)}(K, L)$, we have

$$d^{(m)}(\mathbf{j}, L) \leq d^{(m)}(\mathbf{j}, K) + d^{(m)}(K, L).$$

Finally, taking $\mathbf{j} \in J$ such that $d^{(m)}(\mathbf{j}, L) = d^{(m)}(J, L)$, we obtain

$$\begin{aligned} d^{(m)}(J, L) &\leq d^{(m)}(\mathbf{j}, K) + d^{(m)}(K, L) \\ &\leq d^{(m)}(J, K) + d^{(m)}(K, L), \end{aligned}$$

as required. ■

So the Euclidean distance from one compact set to another satisfies property (M3) of a metric space, but does not satisfy properties (M1) and (M2). In order to overcome this problem, we introduce the following definition of the distance between two compact sets.

Definition 2.10 *Hausdorff distance*

Let $J, K \in \mathcal{K}(\mathbb{R}^m)$, for some $m \in \mathbb{N}$. The **Hausdorff distance** d_{H} between the sets J and K is

$$d_{\mathrm{H}}(J, K) = \max\{d^{(m)}(J, K), d^{(m)}(K, J)\}.$$

Strictly, we should write $d_{\mathrm{H}}^{(m)}$ to indicate the underlying Euclidean space but it is always clear from the context, so we don't bother.

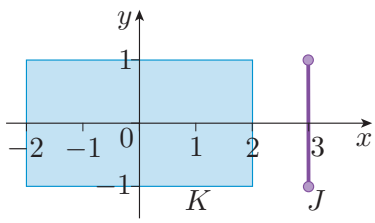

Figure 2.5

For example, consider the sets

$$J = \{(3, y) : |y| \leq 1\} \quad \text{and} \quad K = \{(x, y) : |x| \leq 2 \text{ and } |y| \leq 1\}$$

(see Figure 2.5). In Worked Exercise 2.8 we showed that $d^{(2)}(J, K) = 1$ and $d^{(2)}(K, J) = 5$. So $d_{\mathrm{H}}(J, K) = d_{\mathrm{H}}(K, J) = \max\{1, 5\} = 5$.

Exercise 2.6

Let

$$J = \{(x, y) : x^2 + y^2 \leq 1\} \quad \text{and} \quad K = \{(1, 0)\}.$$

Determine $d_{\mathrm{H}}(J, K)$.

The definition of Hausdorff distance clearly satisfies property (M2) of a metric space, and we now show that it is in fact a metric on $\mathcal{K}(\mathbb{R}^m)$.

Theorem 2.11

The Hausdorff distance, d_{H}, is a metric on $\mathcal{K}(\mathbb{R}^m)$.

Proof Recall that, if J and K belong to $\mathcal{K}(\mathbb{R}^m)$, then the Hausdorff distance between J and K is

$$d_{\mathrm{H}}(J,K) = \max\{d^{(m)}(J,K), d^{(m)}(K,J)\}.$$

We show that d_{H} is a metric on $\mathcal{K}(\mathbb{R}^m)$ by showing that d_{H} has the metric properties (M1)–(M3).

(M1) Let $J, K \in \mathcal{K}(\mathbb{R}^m)$. By the definition of d_{H} it follows that $d_{\mathrm{H}}(J,K) \geq 0$ and $d_{\mathrm{H}}(K,K) = 0$.

It remains only to show that if $d_{\mathrm{H}}(J,K) = 0$, then $J = K$.

If $d_{\mathrm{H}}(J,K) = 0$, then both $d^{(m)}(J,K) = 0$ and $d^{(m)}(K,J) = 0$. Thus,

$$d^{(m)}(\mathbf{j}, K) = 0 \quad \text{for each } \mathbf{j} \in J,$$

and hence $\mathbf{j} \in K$, by Lemma 2.5. So $J \subseteq K$.

Similarly, $K \subseteq J$, and so $J = K$ as required.

Thus (M1) is satisfied.

(M2) Let $J, K \in \mathcal{K}(\mathbb{R}^m)$. By the definition of d_{H},

$$d_{\mathrm{H}}(J,K) = d_{\mathrm{H}}(K,J),$$

and so (M2) is satisfied.

(M3) Let $J, K, L \in \mathcal{K}(\mathbb{R}^m)$. By Lemma 2.9,

$$d^{(m)}(J,L) \leq d^{(m)}(J,K) + d^{(m)}(K,L)$$

and

$$d^{(m)}(L,J) \leq d^{(m)}(L,K) + d^{(m)}(K,J).$$

Since $\max\{a+b, c+d\} \leq \max\{a,d\} + \max\{b,c\}$, for real numbers a, b, c, d, we obtain

$$\begin{aligned} d_{\mathrm{H}}(J,L) &= \max\{d^{(m)}(J,L), d^{(m)}(L,J)\} \\ &\leq \max\{d^{(m)}(J,K) + d^{(m)}(K,L), d^{(m)}(L,K) + d^{(m)}(K,J)\} \\ &\leq \max\{d^{(m)}(J,K), d^{(m)}(K,J)\} + \max\{d^{(m)}(K,L), d^{(m)}(L,K)\} \\ &= d_{\mathrm{H}}(J,K) + d_{\mathrm{H}}(K,L). \end{aligned}$$

Hence (M3) is satisfied.

Since (M1)–(M3) are satisfied, d_{H} is a metric on $\mathcal{K}(\mathbb{R}^m)$. ■

We now have a metric space $(\mathcal{K}(\mathbb{R}^m), d_{\mathrm{H}})$. Before we can apply the Contraction Mapping Theorem, we need the following result.

Theorem 2.12

The metric space $(\mathcal{K}(\mathbb{R}^m), d_{\mathrm{H}})$ is complete.

Recall that a metric space is complete if each Cauchy sequence is convergent. The proof of this theorem is long and complicated. We give it in Subsection 2.4 – for completeness!

2.2 Constructing contraction mappings

We now have a complete metric space $(\mathcal{K}(\mathbb{R}^m), d_{\mathrm{H}})$ to which we can apply the Contraction Mapping Theorem. Our aim is to see how the methods we used in Section 1 to obtain examples of fractals can be described in terms of mappings on $\mathcal{K}(\mathbb{R}^m)$.

In Section 1, we took families of similarities $S_1, S_2, \ldots, S_k$ defined on $\mathbb{R}$ or $\mathbb{R}^2$, and repeatedly applied them to a given compact set K_0. So if

$$K_1 = \bigcup_{i=1}^{k} S_i(K_0)$$

and

$$K_n = \bigcup_{i=1}^{k} S_i(K_{n-1}) \quad \text{for } n \geq 2,$$

then our fractal set is the 'limit' of the sequence of compact sets (K_n). Now that we have a metric for the collections of compact subsets of $\mathbb{R}$ and $\mathbb{R}^2$, we know what the limit of a sequence of compact sets is – it is a compact set J for which

$$d_{\mathrm{H}}(K_n, J) \to 0 \text{ as } n \to \infty.$$

It remains for us to show that *there is* a limit for a sequence of sets K_n determined by similarities $S_1, \ldots, S_k$, and that this limit does not depend on our choice of starting set K_0. We use the Contraction Mapping Theorem for this.

In order to apply the Contraction Mapping Theorem, we describe the construction process in terms of a single mapping from $\mathcal{K}(\mathbb{R})$ to $\mathcal{K}(\mathbb{R})$. So let us define the mapping $S\colon \mathcal{K}(\mathbb{R}) \to \mathcal{K}(\mathbb{R})$ by

$$S(K) = S_1(K) \cup S_2(K) \cup \cdots \cup S_k(K) \quad \text{for each } K \in \mathcal{K}(\mathbb{R}).$$

Then

$$K_n = S(K_{n-1}) \quad \text{for each } n \in \mathbb{N}.$$

Exercise 2.7

Let $S_1, S_2\colon \mathbb{R} \to \mathbb{R}$ be defined by

$$S_1(x) = \tfrac{1}{3}x \quad \text{and} \quad S_2(x) = \tfrac{1}{3}(x+2).$$

Let S be the mapping on $\mathcal{K}(\mathbb{R})$ defined by

$$S(K) = S_1(K) \cup S_2(K),$$

for each $K \in \mathcal{K}(\mathbb{R})$. Find the image of each of the following sets under S, and state whether the image belongs to $\mathcal{K}(\mathbb{R})$.

(a) $[0, 2]$ (b) $[-2, 2]$

In the above exercise, you should have found that each set is mapped by S to a set in $\mathcal{K}(\mathbb{R})$. In fact, *any* set in $\mathcal{K}(\mathbb{R})$ is mapped by S to another set in $\mathcal{K}(\mathbb{R})$, since both S_1 and S_2 are continuous and so map compact sets to

Corollary 5.13 of Chapter 22.

compact sets. Further, the union of two compact sets is also compact, and so S is indeed a mapping from $\mathcal{K}(\mathbb{R})$ to itself.

Proposition 6.4 of Chapter 22.

Similar arguments can be used to prove the following more general result.

Proposition 2.13

Let $S_1, S_2, \ldots, S_k \colon \mathbb{R}^m \to \mathbb{R}^m$ be $(d^{(m)}, d^{(m)})$-continuous, and let S be defined by

$$S(K) = \bigcup_{i=1}^{k} S_i(K) \quad \text{for each } K \in \mathcal{K}(\mathbb{R}^m).$$

Then S maps $\mathcal{K}(\mathbb{R}^m)$ to $\mathcal{K}(\mathbb{R}^m)$.

Remark

Note that each of the fractals in Section 1 is obtained by repeatedly applying mappings of the form described in Proposition 2.13.

Finally, in order to apply the Contraction Mapping Theorem, we require a mapping from $\mathcal{K}(\mathbb{R}^m)$ to itself that is a *contraction mapping*. Note that any similarity with similarity ratio $c < 1$ is a contraction mapping on $\mathbb{R}^m$ with contraction ratio c. In particular, all the similarities that are used to define fractals in Section 1 are contraction mappings.

See Definition 2.1.

In general, suppose that $S_1, S_2, \ldots, S_k \colon \mathbb{R}^m \to \mathbb{R}^m$ are continuous functions and consider the mapping $S \colon \mathcal{K}(\mathbb{R}^m) \to \mathcal{K}(\mathbb{R}^m)$ defined by

$$S(K) = \bigcup_{i=1}^{k} S_i(K) \quad \text{for each } K \in \mathcal{K}(\mathbb{R}^m).$$

We now show that if each of the functions $S_1, S_2, \ldots, S_k$ is a contraction mapping on $\mathbb{R}^m$, then the function S is also a contraction mapping on $\mathcal{K}(\mathbb{R}^m)$.

Proposition 2.14

Let $S_1, S_2, \ldots, S_k \colon \mathbb{R}^m \to \mathbb{R}^m$ be contraction mappings. Then the mapping $S \colon \mathcal{K}(\mathbb{R}^m) \to \mathcal{K}(\mathbb{R}^m)$ defined by

$$S(K) = \bigcup_{i=1}^{k} S_i(K) \quad \text{for each } K \in \mathcal{K}(\mathbb{R}^m)$$

is a contraction mapping with respect to the Hausdorff distance d_{H} on $\mathcal{K}(\mathbb{R}^m)$; that is, there is a number λ such that $0 \leq \lambda < 1$ and

$$d_{\mathrm{H}}(S(J), S(K)) \leq \lambda d_{\mathrm{H}}(J, K),$$

for all $J, K \in \mathcal{K}(\mathbb{R}^m)$.

Proof Each of the mappings $S_1, S_2, \ldots, S_k$ is a contraction mapping, and so there are $0 \le \lambda_1, \lambda_2, \ldots, \lambda_k < 1$ such that

$$d^{(m)}(S_i(\mathbf{x}), S_i(\mathbf{y})) \le \lambda_i d^{(m)}(\mathbf{x}, \mathbf{y}), \quad 1 \le i \le k,$$

for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^m$.

Now let $J, K \in \mathcal{K}(\mathbb{R}^m)$. We show that, for $1 \le i \le k$,

$$d^{(m)}(S_i(J), S_i(K)) \le \lambda_i d^{(m)}(J, K). \tag{1}$$

Indeed, for each $\mathbf{j} \in J$,

$$\begin{aligned} d^{(m)}(S_i(\mathbf{j}), S_i(K)) &= \min_{\mathbf{k}\in K} d^{(m)}(S_i(\mathbf{j}), S_i(\mathbf{k})) \\ &\le \min_{\mathbf{k}\in K} \lambda_i d^{(m)}(\mathbf{j}, \mathbf{k}) = \lambda_i d^{(m)}(\mathbf{j}, K). \end{aligned}$$

So

$$\begin{aligned} d^{(m)}(S_i(J), S_i(K)) &= \max_{\mathbf{j}\in J} d^{(m)}(S_i(\mathbf{j}), S_i(K)) \\ &\le \max_{\mathbf{j}\in J} \lambda_i d^{(m)}(\mathbf{j}, K) = \lambda_i d^{(m)}(J, K). \end{aligned}$$

Thus (1) is true.

Now

$$\begin{aligned} d^{(m)}(S(J), S(K)) &= d^{(m)}\left(\bigcup_{i=1}^{k} S_i(J), \bigcup_{j=1}^{k} S_j(K)\right) \\ &= \max_{\mathbf{x}\in\bigcup_{i=1}^{k} S_i(J)} d^{(m)}\left(\mathbf{x}, \bigcup_{j=1}^{k} S_j(K)\right) \\ &= \max_{1\le i\le k} \max_{\mathbf{x}\in S_i(J)} d^{(m)}\left(\mathbf{x}, \bigcup_{j=1}^{k} S_j(K)\right) \\ &= \max_{1\le i\le k} d^{(m)}\left(S_i(J), \bigcup_{j=1}^{k} S_j(K)\right). \end{aligned}$$

Since $S_i(K) \subseteq \bigcup_{j=1}^{k} S_j(K)$ for $i = 1, 2, \ldots, k$, the result of Exercise 2.5(b) implies that

$$d^{(m)}(S(J), S(K)) \le \max_{1\le i\le k} d^{(m)}(S_i(J), S_i(K)).$$

This estimate, together with inequality (1), implies that

$$d^{(m)}(S(J), S(K)) \le \max_{1\le i\le k} \lambda_i d^{(m)}(J, K) = \left(\max_{1\le i\le k} \lambda_i\right) d^{(m)}(J, K).$$

The same arguments show that

$$d^{(m)}(S(K), S(J)) \le \max_{1\le i\le k} \lambda_i d^{(m)}(K, J) = \left(\max_{1\le i\le k} \lambda_i\right) d^{(m)}(K, J),$$

and so

$$d_{\mathrm{H}}(S(J), S(K)) \le \left(\max_{1\le i\le k} \lambda_i\right) d_{\mathrm{H}}(J, K).$$

But $0 \le \max_{1\le i\le k} \lambda_i < 1$, and so S is a contraction mapping on $\mathcal{K}(\mathbb{R}^m)$ with contraction ratio $\lambda = \max_{1\le i\le k} \lambda_i$. ■

2.3 Iterated function schemes

We have shown that, for a given family of contraction mappings $S_1, S_2, \ldots, S_k \colon \mathbb{R}^m \to \mathbb{R}^m$, the mapping $S \colon \mathcal{K}(\mathbb{R}^m) \to \mathcal{K}(\mathbb{R}^m)$ defined by

$$S(K) = \bigcup_{i=1}^{k} S_i(K) \quad \text{for each } K \in \mathcal{K}(\mathbb{R}^m),$$

is a contraction mapping with respect to the Hausdorff distance d_{H} on $\mathcal{K}(\mathbb{R}^m)$. Since $(\mathcal{K}(\mathbb{R}^m), d_{\mathrm{H}})$ is a complete metric space, we can apply the Contraction Mapping Theorem to obtain the following important result.

Theorem 2.15

Let $S_1, S_2, \ldots, S_k \colon \mathbb{R}^m \to \mathbb{R}^m$ be contraction mappings, and let the mapping $S \colon \mathcal{K}(\mathbb{R}^m) \to \mathcal{K}(\mathbb{R}^m)$ be defined by

$$S(K) = \bigcup_{i=1}^{k} S_i(K) \quad \text{for each } K \in \mathcal{K}(\mathbb{R}^m).$$

Then there is a unique set $K_S \in \mathcal{K}(\mathbb{R}^m)$ such that

$$S(K_S) = K_S.$$

Moreover, for each $K \in \mathcal{K}(\mathbb{R}^m)$, the sequence of iterates

$$K,\ S(K),\ S^2(K),\ S^3(K),\ \ldots$$

converges to K_S; that is,

$$d_{\mathrm{H}}(S^n(K), K_S) \to 0 \text{ as } n \to \infty.$$

This is our main result.

Remark

This result is important because for many families of contraction mappings $S_1, S_2, \ldots, S_k$, the set K_S is a fractal. In fact, this result gives a simple method of constructing these fractals – we take a compact set and repeatedly apply the mapping S. This method of generating a fractal is called an *iterated function scheme*.

Since

$$S(K_S) = \bigcup_{i=1}^{k} S_i(K_S) = K_S,$$

we say that the set K_S is *invariant* for the contraction mappings $S_1, S_2, \ldots, S_k$.

Definition 2.16 *Invariant*

A set A is **invariant** for the mappings $S_1, S_2, \ldots, S_k \colon \mathbb{R}^m \to \mathbb{R}^m$ if

$$\bigcup_{i=1}^{k} S_i(A) = A.$$

Remark

A self-similar set is just a set that is invariant for *similarities* $S_1, S_2, \ldots, S_k \colon \mathbb{R}^m \to \mathbb{R}^m$.

We now have a rigorous justification of the process that we used in Section 1 to generate fractals.

The von Koch curve

To construct the von Koch curve, we take the compact set $K_0 = [0, 1]$ and repeatedly apply the mapping $S \colon \mathcal{K}(\mathbb{R}^2) \to \mathcal{K}(\mathbb{R}^2)$ given by

$$S(K) = S_1(K) \cup S_2(K) \cup S_3(K) \cup S_4(K) \quad \text{for each } K \in \mathcal{K}(\mathbb{R}^2)$$

(Figure 2.6), where the similarities $S_1, S_2, S_3, S_4 \colon \mathbb{R}^2 \to \mathbb{R}^2$ are as follows.

S_1 scales by a factor of $\frac{1}{3}$.

S_2 scales by a factor of $\frac{1}{3}$, rotates anticlockwise through $\frac{\pi}{3}$ about the point $(0, 0)$, and translates by $\left(\frac{1}{3}, 0\right)$.

S_3 scales by a factor of $\frac{1}{3}$, rotates clockwise by $\frac{\pi}{3}$ about the point $(0, 0)$, and translates by $(\frac{1}{2}, \frac{\sqrt{3}}{6})$.

S_4 scales by a factor of $\frac{1}{3}$ and translates by $\left(\frac{2}{3}, 0\right)$.

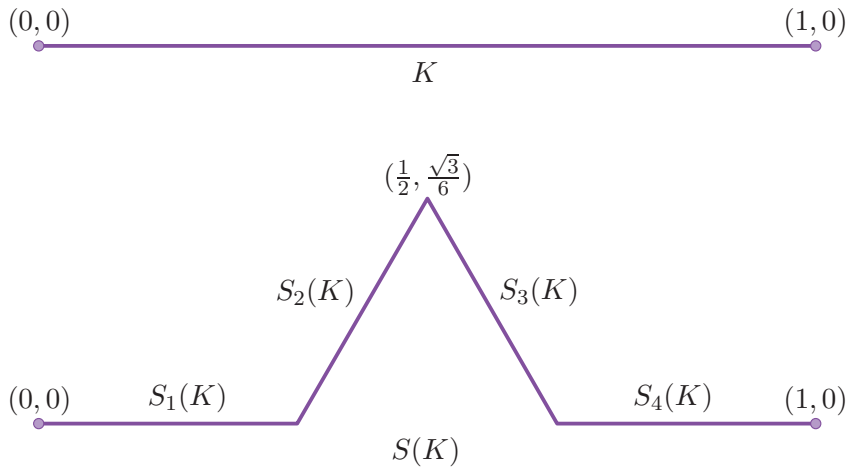

Figure 2.6

We define the von Koch curve to be the limit of the curves $K_n = S^n(K_0)$. In Section 1, we did not justify the existence of such a limit. Since each of

the mappings S_1, S_2, S_3, S_4 is a contraction mapping (each with contraction ratio $\frac{1}{3}$), it now follows from Theorem 2.15 that such a limit does indeed exist.

We also claimed in Section 1 that the von Koch curve K_V satisfies the following equation, and is therefore self-similar:

$$K_V = S_1(K_V) \cup S_2(K_V) \cup S_3(K_V) \cup S_4(K_V) = S(K_V),$$

that is, K_V is invariant for S_1, S_2, S_3, S_4. This also follows from Theorem 2.15.

Many other fractal curves can be constructed similarly, as Exercise 2.8 illustrates.

Exercise 2.8

Let $S_1, S_2, S_3, S_4, S_5 \colon \mathbb{R}^2 \to \mathbb{R}^2$ be similarities defined by

$$S_1(x,y) = \left(\tfrac{1}{3}x, \tfrac{1}{3}y\right), \qquad S_2(x,y) = \left(\tfrac{1}{3}(y+1), \tfrac{1}{3}x\right),$$
$$S_3(x,y) = \left(\tfrac{1}{3}(x+1), \tfrac{1}{3}(y+1)\right), \qquad S_4(x,y) = \left(\tfrac{1}{3}(-y+2), \tfrac{1}{3}x\right),$$
$$S_5(x,y) = \left(\tfrac{1}{3}(x+2), \tfrac{1}{3}y\right).$$

Let

$$S(K) = \bigcup_{i=1}^{5} S_i(K) \quad \text{for each } K \in \mathcal{K}(\mathbb{R}^2),$$

and let $K_0 = \{(x,y) : 0 \le x \le 1 \text{ and } y = 0\}$.

(a) Sketch the sets $S(K_0)$ and $S^2(K_0)$.

(b) Determine the length of $S^n(K_0)$, for each $n \in \mathbb{N}$.

(c) Determine whether $\lim_{n\to\infty} S^n(K_0)$ exists.

Cantor sets

In Section 1, our first example of a fractal was the middle-third Cantor set. To construct it, we take the compact set $K_0 = [0,1]$ and repeatedly apply the mapping $S \colon \mathcal{K}(\mathbb{R}) \to \mathcal{K}(\mathbb{R})$ defined by

$$S(K) = S_1(K) \cup S_2(K),$$

where

$$S_1(x) = \tfrac{1}{3}x \quad \text{and} \quad S_2(x) = \tfrac{1}{3}(x+2).$$

We define the middle-third Cantor set to be the *intersection* of the nested sequence of sets $K_n = S^n(K_0)$. Since each of the mappings S_1 and S_2 is a contraction mapping (with contraction ratio $\frac{1}{3}$), it follows from Theorem 2.15 that the sets K_n converge to a unique limit set K_S (with respect to Hausdorff distance) as $n \to \infty$. It seems reasonable to expect that this limit set is the same as K_C, the middle-third Cantor set. The following lemma shows this to be the case.

Lemma 2.17

Let $S_1, S_2, \ldots, S_k \colon \mathbb{R}^m \to \mathbb{R}^m$ be contraction mappings, and let $S \colon \mathcal{K}(\mathbb{R}^m) \to \mathcal{K}(\mathbb{R}^m)$ be the mapping defined by

$$S(K) = \bigcup_{i=1}^{k} S_i(K) \quad \text{for each } K \in \mathcal{K}(\mathbb{R}^m).$$

Let $S(K_0) \subseteq K_0$, for some $K_0 \in \mathcal{K}(\mathbb{R}^m)$; then

$$\bigcap_{n \in \mathbb{N}} S^n(K_0) = K_S,$$

where K_S is the unique set in $\mathcal{K}(\mathbb{R}^m)$ such that $S(K_S) = K_S$.

Proof For simplicity, let $K_n = S^n(K_0)$ and $J = \bigcap_{n \in \mathbb{N}} K_n$. Then, for each $n \in \mathbb{N}$,

$$K_{n+1} = S^{n+1}(K_0) = S^n(S(K_0)) \subseteq S^n(K_0) = K_n,$$

and so the sets K_n are nested.

To show that $J = K_S$, we show that both $J - K_S$ and $K_S - J$ are empty, and so $J \subseteq K_S$ and $K_S \subseteq J$.

We use a proof by contradiction.

First, suppose that $J - K_S \neq \emptyset$, and so there is an $\mathbf{x} \in J - K_S$.

Then by Lemma 2.5 there is an $\varepsilon > 0$ such that

$$d^{(m)}(\mathbf{x}, K_S) = \varepsilon.$$

Since $\mathbf{x} \in J$, it follows that $\mathbf{x} \in K_n$, for each $n \in \mathbb{N}$, and so

$$d_{\mathrm{H}}(K_n, K_S) \geq d^{(m)}(K_n, K_S) \geq d^{(m)}(\mathbf{x}, K_S) = \varepsilon.$$

This contradicts Theorem 2.15, and so there is no $\mathbf{x} \in J - K_S$; thus $J \subseteq K_S$.

Now suppose that $K_S - J \neq \emptyset$, and so there is an $\mathbf{x} \in K_S - J$.

Since $\mathbf{x} \notin J$, there is an $N \in \mathbb{N}$ with $\mathbf{x} \notin K_N$. Hence by Lemma 2.5 there is an $\varepsilon > 0$ such that

$$d^{(m)}(\mathbf{x}, K_N) = \varepsilon.$$

Since the sets K_n are nested, we find that, for each $n \geq N$,

$$\begin{aligned} d_{\mathrm{H}}(K_n, K_S) \geq d^{(m)}(K_S, K_n) &\geq d^{(m)}(\mathbf{x}, K_n) \\ &\geq d^{(m)}(\mathbf{x}, K_N) \quad \text{by Lemma 2.6,} \\ &= \varepsilon. \end{aligned}$$

This contradicts Theorem 2.15, and so there is no $\mathbf{x} \in K_S - J$; thus $K_S \subseteq J$.

Hence $K_S = J$, as required. ■

So, as we expected, the middle-third Cantor set $K_{\mathrm{C}} = \bigcap_{n \in \mathbb{N}} K_n$ is the limit of the sets $K_n = S^n(K_0)$. It follows from Theorem 2.15 that the

middle-third Cantor set K_C satisfies the following equation, and is therefore self-similar:

$$K_C = S_1(K_C) \cup S_2(K_C) = S(K_C). \tag{2}$$

In fact, Theorem 2.15 tells us that the middle-third Cantor set is the *only* set in $\mathcal{K}(\mathbb{R})$ that is invariant for S_1 and S_2. This means that the middle-third Cantor set can be *defined* as the unique non-empty compact solution to (2). So this intricate set can be defined by a simple equation.

Note that there are solutions to (2) that are *not* non-empty and compact, as the following exercise illustrates.

Exercise 2.9

Show that for K equal to any of the sets $\emptyset$, $\mathbb{R}$ and $[0, \infty)$,

$$K = S_1(K) \cup S_2(K),$$

where $S_1, S_2 \colon \mathbb{R} \to \mathbb{R}$ are given by $S_1(x) = \frac{1}{3}x$ and $S_2(x) = \frac{1}{3}(x + 2)$.

It follows from Theorem 2.15 that if we take K_0 to be *any* non-empty compact subset of $\mathbb{R}$, then the middle-third Cantor set is the limit of the sets $K_n = S^n(K_0)$. Further, it follows from Lemma 2.17 that if $S(K_0) \subseteq K_0$, then the middle-third Cantor set is the intersection of the sets K_n.

Exercise 2.10

Let K_C be the middle-third Cantor set, and let $S \colon \mathcal{K}(\mathbb{R}) \to \mathcal{K}(\mathbb{R})$ be the mapping defined by

$$S(K) = S_1(K) \cup S_2(K),$$

where $S_1, S_2 \colon \mathbb{R} \to \mathbb{R}$ are given by

$$S_1(x) = \tfrac{1}{3}x \quad \text{and} \quad S_2(x) = \tfrac{1}{3}(x + 2).$$

Show that

$$\bigcap_{k=0}^{\infty} S^k([0, 2]) = K_C.$$

Note that if K_0 is *not* compact, then the sets $K_n = S^n(K_0)$ do not necessarily converge to the middle-third Cantor set. For example, if $K_0 = [0, \infty)$, then

$$S(K_0) = S_1(K_0) \cup S_2(K_0) = [0, \infty) \cup [\tfrac{2}{3}, \infty) = [0, \infty) = K_0,$$

and so $S^n(K_0) = K_0$, for each $n \in \mathbb{N}$.

The Sierpiński gasket

Our other main example of a fractal is the Sierpiński gasket. To construct it, we take K_0 to be the equilateral triangle whose base is the line segment from $(0,0)$ to $(1,0)$, and repeatedly apply the mapping $S\colon \mathcal{K}(\mathbb{R}^2) \to \mathcal{K}(\mathbb{R}^2)$ given by

$$S(K) = S_1(K) \cup S_2(K) \cup S_3(K),$$

where the mappings $S_1, S_2, S_3\colon \mathbb{R}^2 \to \mathbb{R}^2$ are defined by

$$S_1(x,y) = \left(\tfrac{1}{2}x, \tfrac{1}{2}y\right), \qquad S_2(x,y) = \left(\tfrac{1}{2}(x+1), \tfrac{1}{2}y\right),$$

$$S_3(x,y) = \left(\tfrac{1}{4}(2x+1), \tfrac{1}{4}(2y+\sqrt{3})\right).$$

We define the Sierpiński gasket to be the *intersection* of the nested sequence of sets $K_n = S^n(K_0)$. Since each of the mappings S_1, S_2 and S_3 is a contraction mapping (with contraction ratio $\frac{1}{2}$), it follows from Theorem 2.15 that the sets K_n converge to a unique limit set (with respect to Hausdorff distance) as $n \to \infty$. Since $S(K_0) \subseteq K_0$, it follows from Lemma 2.17 that this limit set is the Sierpiński gasket.

We now know from Theorem 2.15 that the Sierpiński gasket can be obtained by repeatedly applying the mapping S to *any* non-empty compact subset of $\mathbb{R}^2$. Further, it follows from Theorem 2.15 that the Sierpiński gasket can be defined to be the unique non-empty compact set that is invariant for S_1, S_2 and S_3.

In the next exercise, we ask you to show how similarities can also be used to define simple sets.

Exercise 2.11

Let $S\colon \mathcal{K}(\mathbb{R}^2) \to \mathcal{K}(\mathbb{R}^2)$ be defined by

$$S(K) = S_1(K) \cup S_2(K),$$

where $S_1, S_2\colon \mathbb{R}^2 \to \mathbb{R}^2$ are given by

$$S_1(x,y) = \left(\tfrac{1}{2}x, \tfrac{1}{2}y\right) \quad \text{and} \quad S_2(x,y) = \left(\tfrac{1}{2}x + \tfrac{1}{2}, \tfrac{1}{2}y + \tfrac{1}{2}\right).$$

(a) Sketch the sets $K_0 = \{(x,y) : 0 \leq x \leq 1 \text{ and } 0 \leq y \leq 1\}$ and $K_1 = S(K_0)$.

(b) Let $K_n = S^n(K_0)$, for each $n \in \mathbb{N}$. Verify that

$$\lim_{n\to\infty} K_n = \{(x,y) : y = x \text{ and } 0 \leq x \leq 1\}.$$

2.4 Proof of Theorem 2.12

This proof will not be assessed.

See Definition 2.1 of Chapter 23.

We now prove that $(\mathcal{K}(\mathbb{R}^m), d_{\mathrm{H}})$ is a complete metric space. Recall that a metric space (X,d) is complete if each Cauchy sequence in X converges to a point $x \in X$.

Let (K_n) be a Cauchy sequence in $\mathcal{K}(\mathbb{R}^m)$. Then K_n is a non-empty compact set for each $n \in \mathbb{N}$, and for each $\varepsilon > 0$, there is an $N \in \mathbb{N}$ such that

$$d_{\mathrm{H}}(K_n, K_{n'}) < \varepsilon \quad \text{for all } n, n' > N.$$

In order to show that $(\mathcal{K}(\mathbb{R}^m), d_{\mathrm{H}})$ is complete, we find a non-empty compact set J with $d_{\mathrm{H}}(K_n, J) \to 0$ as $n \to \infty$. Since (K_n) is a Cauchy sequence, it is enough to find a convergent subsequence of (K_n).

See Lemma 1.6 of Chapter 23.

Since (K_n) is a Cauchy sequence, for each $i \in \mathbb{N} \cup \{0\}$ there is an $N(i) \in \mathbb{N}$ with $N(i+1) > N(i)$ such that

$$d_{\mathrm{H}}(K_n, K_{n'}) < \left(\tfrac{1}{2}\right)^i \quad \text{for all } n, n' > N(i). \tag{3}$$

We show that $(K_{N(i)})_{i \in \mathbb{N}}$ is d_{H}-convergent in $\mathcal{K}(\mathbb{R}^m)$.

We start by finding a candidate for the limit of $K_{N(i)}$ as $i \to \infty$.

We consider the sets

$$J_i = \left\{ \mathbf{x} \in \mathbb{R}^m : d^{(m)}(\mathbf{x}, K_{N(i)}) \le \left(\tfrac{1}{2}\right)^{i-2} \right\} \quad \text{for } i \in \mathbb{N}.$$

If $\mathbf{x} \in K_{N(i)}$, then it follows from Lemma 2.5 that $d^{(m)}(\mathbf{x}, K_{N(i)}) = 0$, and so $\mathbf{x} \in J_i$. Thus, $K_{N(i)} \subseteq J_i$ and hence, by Exercise 2.5(a),

$$d^{(m)}(K_{N(i)}, J_i) = 0. \tag{4}$$

Notice also that

$$d^{(m)}(J_i, K_{N(i)}) = \max\{d^{(m)}(\mathbf{x}, K_{N(i)}) : \mathbf{x} \in J_i\} \le \left(\tfrac{1}{2}\right)^{i-2}, \tag{5}$$

and so

$$d_{\mathrm{H}}(J_i, K_{N(i)}) \le \left(\tfrac{1}{2}\right)^{i-2}. \tag{6}$$

Let

$$J = \bigcap_{i \in \mathbb{N}} J_i.$$

This set is our candidate for the limit of the sequence $(K_{N(i)})_{i \in \mathbb{N}}$.

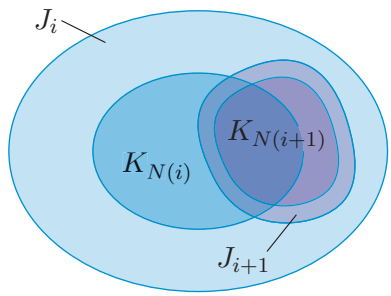

Figure 2.7

We show that

(a) $J \in \mathcal{K}(\mathbb{R}^m)$

(b) $\lim_{i \to \infty} d_{\mathrm{H}}(K_{N(i)}, J) = 0$.

Proof of (a)

We start by noting that each of the sets J_i is a non-empty compact set and so belongs to $\mathcal{K}(\mathbb{R}^m)$.

We ask you to prove this in the exercises.

We now estimate the $d^{(m)}$-distance of J_{i+1} from $K_{N(i)}$. It follows from Lemma 2.9, (5) and (3) that

$$\begin{aligned} d^{(m)}(J_{i+1}, K_{N(i)}) &\le d^{(m)}(J_{i+1}, K_{N(i+1)}) + d^{(m)}(K_{N(i+1)}, K_{N(i)}) \\ &\le \left(\tfrac{1}{2}\right)^{i-1} + d^{(m)}(K_{N(i+1)}, K_{N(i)}) \\ &< \left(\tfrac{1}{2}\right)^{i-1} + \left(\tfrac{1}{2}\right)^{i-1} = \left(\tfrac{1}{2}\right)^{i-2}. \end{aligned}$$

When applying inequality (3), we use the fact that $N(i+1) > N(i) > N(i-1)$.

Hence if $\mathbf{x} \in J_{i+1}$, then $d^{(m)}(\mathbf{x}, K_{N(i)}) < 1/2^{i-2}$; this implies that $\mathbf{x} \in J_i$, and so $J_{i+1} \subseteq J_i$ for each $i \in \mathbb{N}$ (see Figure 2.7). Thus J is the intersection of a nested sequence of non-empty compact sets and is therefore non-empty and compact.

See Theorem 6.3 of Chapter 22.

Proof of (b)

To prove (b) it is enough to show that

$$d_{\mathrm{H}}(J_i, J) \to 0 \text{ as } i \to \infty, \tag{7}$$

since (M3) for d_{H} together with (6) implies that

$$d_{\mathrm{H}}(K_{N(i)}, J) \le d_{\mathrm{H}}(K_{N(i)}, J_i) + d_{\mathrm{H}}(J_i, J) \le \left(\tfrac{1}{2}\right)^{i-2} + d_{\mathrm{H}}(J_i, J),$$

and so if $d_{\mathrm{H}}(J_i, J) \to 0$ as $i \to \infty$, then $d_{\mathrm{H}}(K_{N(i)}, J) \to 0$ as $i \to \infty$.

To show that (7) holds, we need show only that $d^{(m)}(J_i, J) \to 0$ as $i \to \infty$, since $J \subseteq J_i$ for each i implies that $d^{(m)}(J, J_i) = 0$ (by Exercise 2.5(a)).

Thus for each $i \in \mathbb{N}$, we need to find an upper estimate for

$$d^{(m)}(J_i, J) = \max\{d^{(m)}(\mathbf{x}, J) : \mathbf{x} \in J_i\}.$$

So we fix $i \in \mathbb{N}$ and let $\mathbf{x} \in J_i$.

We define a sequence of points $(\mathbf{y}_j)_{j\in\mathbb{N}}$ as follows. Take

$$\mathbf{y}_1 \in J_{i+1} \text{ such that } d^{(m)}(\mathbf{x}, \mathbf{y}_1) \le \left(\tfrac{1}{2}\right)^{i-2},$$
$$\mathbf{y}_2 \in J_{i+2} \text{ such that } d^{(m)}(\mathbf{y}_1, \mathbf{y}_2) \le \left(\tfrac{1}{2}\right)^{i-1},$$

and, more generally, for $j \ge 2$ choose

$$\mathbf{y}_j \in J_{i+j} \text{ such that } d^{(m)}(\mathbf{y}_{j-1}, \mathbf{y}_j) \le \left(\tfrac{1}{2}\right)^{i+j-3}. \tag{8}$$

This is possible since it follows from Lemma 2.9, (5), (3) and (4) that

$$\begin{aligned} d^{(m)}(J_{i+j}, J_{i+j-1}) &\le d^{(m)}(J_{i+j}, K_{N(i+j)}) + d^{(m)}(K_{N(i+j)}, K_{N(i+j-1)}) \\ &\qquad + d^{(m)}(K_{N(i+j-1)}, J_{i+j-1}) \\ &< \left(\tfrac{1}{2}\right)^{i+j-2} + \left(\tfrac{1}{2}\right)^{i+j-2} + 0 = \left(\tfrac{1}{2}\right)^{i+j-3}. \end{aligned}$$

It is easy to check that $(\mathbf{y}_j)$ is a $d^{(m)}$-Cauchy sequence in $\mathbb{R}^m$: for if $n, n' \in \mathbb{N}$ with $n' > n$, then (M3) for $d^{(m)}$ together with (8) gives

$$\begin{aligned} d^{(m)}(\mathbf{y}_{n'}, \mathbf{y}_n) &\le d^{(m)}(\mathbf{y}_{n'}, \mathbf{y}_{n'-1}) + d^{(m)}(\mathbf{y}_{n'-1}, \mathbf{y}_{n'-2}) + \cdots + d^{(m)}(\mathbf{y}_{n+1}, \mathbf{y}_n) \\ &\le \left(\tfrac{1}{2}\right)^{i+n'-3} + \left(\tfrac{1}{2}\right)^{i+n'-4} + \cdots + \left(\tfrac{1}{2}\right)^{i+n-2} \\ &\le \left(\tfrac{1}{2}\right)^{i+n-1} \to 0 \quad \text{as } n \to \infty. \end{aligned}$$

Since $(\mathbb{R}^m, d^{(m)})$ is complete, the sequence has a limit $\mathbf{y} \in \mathbb{R}^m$.

Moreover, since the sequence $(\mathbf{y}_j)$ lies in the compact set J_i, its limit $\mathbf{y}$ is also in J_i. For $M \in \mathbb{N}$, the subsequence $(\mathbf{y}_j)_{j\ge M}$ is a Cauchy sequence in the compact set J_{i+M}, and so the limit $\mathbf{y}$ must also be in J_{i+M}. This is true for each $M \in \mathbb{N}$, and so

$$\mathbf{y} \in \bigcap_{M\in\mathbb{N}} J_{i+M} = J.$$

Hence $d^{(m)}(\mathbf{x}, J) \leq d^{(m)}(\mathbf{x}, \mathbf{y})$. By repeatedly using the Triangle Inequality and (8), we estimate that, for any $N \in \mathbb{N}$,

$$\begin{aligned} d^{(m)}(\mathbf{x}, \mathbf{y}) &\leq d^{(m)}(\mathbf{x}, \mathbf{y}_1) + d^{(m)}(\mathbf{y}_1, \mathbf{y}_2) + \cdots + d^{(m)}(\mathbf{y}_{N-1}, \mathbf{y}_N) + d^{(m)}(\mathbf{y}_N, \mathbf{y}) \\ &\leq \left(\tfrac{1}{2}\right)^{i-2} + \left(\tfrac{1}{2}\right)^{i-1} + \cdots + \left(\tfrac{1}{2}\right)^{i+N-3} + d^{(m)}(\mathbf{y}_N, \mathbf{y}) \\ &\leq \sum_{j \in \mathbb{N}} \left(\tfrac{1}{2}\right)^{i+j-3} + d^{(m)}(\mathbf{y}_N, \mathbf{y}) \\ &= \left(\tfrac{1}{2}\right)^{i-3} + d^{(m)}(\mathbf{y}_N, \mathbf{y}). \end{aligned}$$

However, $d^{(m)}(\mathbf{y}_N, \mathbf{y}) \to 0$ as $N \to \infty$, and so

$$d^{(m)}(\mathbf{x}, \mathbf{y}) \leq \left(\tfrac{1}{2}\right)^{i-3}.$$

This implies that

$$d^{(m)}(\mathbf{x}, J) \leq \left(\tfrac{1}{2}\right)^{i-3}.$$

Since this is true for any point $\mathbf{x} \in J_i$, it follows that

$$d^{(m)}(J_i, J) \leq \left(\tfrac{1}{2}\right)^{i-3},$$

and so $d^{(m)}(J_i, J) \to 0$ as $i \to \infty$, as required.

Since both (a) and (b) hold, $(\mathcal{K}(\mathbb{R}^m), d_{\mathrm{H}})$ is a complete metric space.

3 Dimensions of fractals

In Section 1 you saw several examples of fractals, each of which is hard to describe in terms of classical geometry. For example, the von Koch curve seems to be 'smaller' than two-dimensional sets such as rectangles (since it has zero area) and yet seems to be 'larger' than one-dimensional sets such as line segments (since it is bounded but does not have finite length).

We now introduce a way of defining dimension that allows the dimension of a set to be a non-integer number. This enables us to assign a dimension to the von Koch curve that is greater than 1 but less than 2.

3.1 Box dimension

There are many definitions of dimension that take non-integer values. Here we look at one of the most widely used – the *box dimension*. We begin with an informal discussion of box dimension; for simplicity, we restrict our attention to subsets of the plane.

Suppose that we wish to estimate the dimension of a bounded set $A \subset \mathbb{R}^2$. We begin by covering the plane with a grid made up of boxes of side d. We then count how many boxes in the grid intersect the set A (Figure 3.1).

We adopt the following notation.

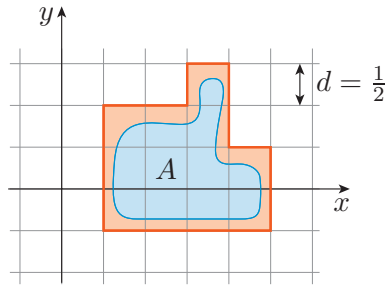

Figure 3.1

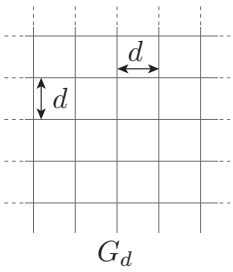

Figure 3.2

Definition 3.1 *Grid of side d*

Let $d \in (0,1)$. The **grid** G_d is a particular set of closed boxes of side d:

$$G_d = \{[nd,(n+1)d] \times [md,(m+1)d] : n,m \in \mathbb{Z}\} \quad \text{(see Figure 3.2)}.$$

The number of (closed) boxes in G_d that intersect a set A is denoted by $N_d(A)$.

For example, Figure 3.1 shows a set A that meets 12 boxes in the grid $G_{1/2}$, so $N_{1/2}(A) = 12$.

Exercise 3.1

This exercise shows that applying a translation to a set A can change the number $N_d(A)$.

Let A be a horizontal line segment of length 1. Calculate $N_{1/4}(A)$ when

(a) $A = \{(x,y) : 0 \le x \le 1 \text{ and } y = 0\}$

(b) $A = \{(x,y) : 0 \le x \le 1 \text{ and } y = \frac{1}{8}\}$

(c) $A = \{(x,y) : -\frac{1}{8} \le x \le \frac{7}{8} \text{ and } y = \frac{1}{8}\}$

(d) $A = \{(x,y) : -\frac{1}{8} \le x \le \frac{7}{8} \text{ and } y = 0\}$.

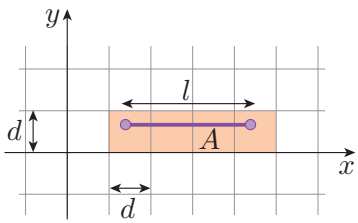

Figure 3.3

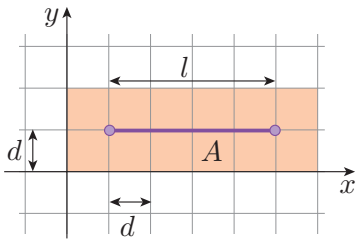

Figure 3.4

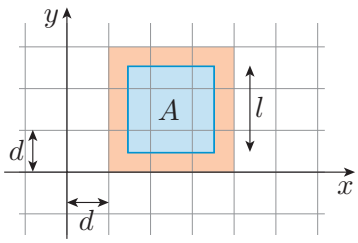

Figure 3.5

Let A be a horizontal line segment of length $l > 0$. For $d > 0$, whatever the position of A relative to the grid G_d, we have (Figure 3.3)

$$N_d(A) \ge \frac{l}{d},$$

since

$$N_d(A) \times (\text{the side length of each box in } G_d) \ge \text{the length of } A.$$

The number $N_d(A)$ is highest when A is on one of the grid lines and has its endpoints at vertices of the grid (as in Exercise 3.1(a); see Figure 3.4).

In this case, $N_d(A) = 2\left(\frac{l}{d} + 2\right)$. So, if $d < l$, then $1 < \frac{l}{d}$, so $2 < 2\frac{l}{d}$, so

$$N_d(A) \le 2\left(\frac{l}{d} + 2 \times \frac{l}{d}\right) = \frac{6l}{d},$$

so

$$\frac{l}{d} \le N_d(A) \le \frac{6l}{d}.$$

Thus, although changing the position of the straight line A relative to the grid can change the value of $N_d(A)$, it can change it by a factor of at most 6. Whatever the position of A, the number $N_d(A)$ 'grows like' $1/d$ for $d < l$.

We now consider the case when A is a square of side $l > 0$ with sides parallel to the axes. Whatever the position of A, we have (Figure 3.5)

$$N_d(A) \ge \left(\frac{l}{d}\right)^2,$$

since

$$N_d(A) \times (\text{the area of each box in } G_d) \geq \text{the area of } A.$$

The number $N_d(A)$ is highest when the edges of A are on the grid lines and have their endpoints at vertices of the grid (Figure 3.6).

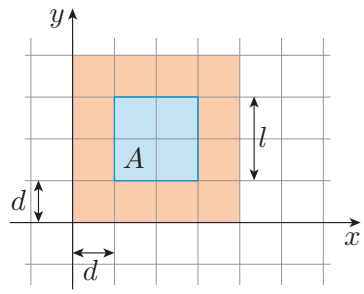

Figure 3.6

In this case,

$$N_d(A) = \left(\frac{l}{d} + 2\right) \times \left(\frac{l}{d} + 2\right).$$

So, if $d < l$, then

$$N_d(A) \leq \left(\frac{l}{d} + \left(2 \times \frac{l}{d}\right)\right) \times \left(\frac{l}{d} + \left(2 \times \frac{l}{d}\right)\right) = \frac{3l}{d} \times \frac{3l}{d} = \left(\frac{3l}{d}\right)^2,$$

so

$$\frac{l^2}{d^2} \leq N_d(A) \leq \frac{9l^2}{d^2}.$$

Hence, whatever the position of the square A relative to the grid, the number $N_d(A)$ 'grows like' $1/d^2$ for $d < l$.

For both these examples, provided d is sufficiently small, $N_d(A)$ 'grows like' $1/d^s$, where s is our intuitive idea of the dimension of the set A. This suggests that if we have *any* bounded set A for which $N_d(A)$ 'grows like' $1/d^s$, then we can *define* its dimension to be s. This gives a definition of dimension that can take non-integer values, but that agrees with our existing ideas of dimension for simple sets such as line segments and squares.

We now formalise this idea. Suppose that there are numbers m and M such that $0 < m \leq M < \infty$ and

$$\frac{m}{d^s} \leq N_d(A) \leq \frac{M}{d^s} \quad \text{for } 0 < d < 1.$$

Taking logarithms, we obtain

$$\log m - s \log d \leq \log N_d(A) \leq \log M - s \log d.$$

Dividing by $-\log d$, we obtain

$$\frac{\log m}{-\log d} + s \leq \frac{\log N_d(A)}{-\log d} \leq \frac{\log M}{-\log d} + s. \qquad (1)$$

Here $\log d$ is negative, since $0 < d < 1$, so $-\log d$ is positive.

Finally, we take the limit as $d \to 0$. As $d \to 0$, we have $\log d \to -\infty$ and thus $-\log d \to +\infty$. Since m and M are constant, $\frac{\log m}{-\log d}$ and $\frac{\log M}{-\log d}$ both tend to zero as $d \to 0$. Hence, in the limit, inequality (1) becomes

$$s \leq \lim_{d \to 0} \frac{\log N_d(A)}{-\log d} \leq s,$$

and thus

$$s = \lim_{d \to 0} \frac{\log N_d(A)}{-\log d}.$$

We now give a formal definition of box dimension.

Box dimension is sometimes referred to as *box-counting dimension*.

Definition 3.2 *Box dimension*

Let A be a bounded subset of $\mathbb{R}^2$. The **box dimension** of A is

$$\dim A = \lim_{d \to 0} \frac{\log N_d(A)}{-\log d},$$

provided that this limit exists.

See Definition 3.1.

Recall that $N_d(A)$ is the number of (closed) boxes in the grid G_d that intersect A. We have seen that $N_d(A)$ can change when A is subjected to a translation; however, the *limit* is well defined, and thus our definition of box dimension makes sense.

Worked Exercise 3.3

Determine the box dimension of the middle-third Cantor set K_C.

Solution

We regard K_C as a subset of the plane by embedding it in the x-axis.

Our objective is to find

$$\lim_{d \to 0} \frac{\log N_d(K_C)}{-\log d}.$$

For a given $d > 0$, we estimate $N_d(K_C)$ by comparing it with $N_d(K_n)$, where n is chosen so that 3^{-n} is approximately equal to d; Figure 3.7 shows $G_{3^{-2}}$ and K_2.

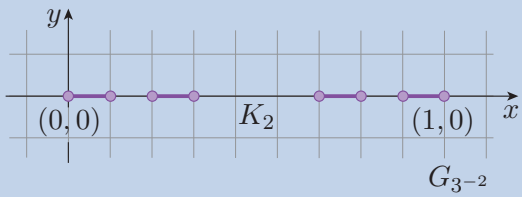

Figure 3.7

We note first that, at the nth stage of the construction of K_C, the set K_n consists of 2^n line segments each of length $\left(\frac{1}{3}\right)^n$. The set K_n contains K_C, and since the endpoints of the intervals making up K_n are not removed in the construction, each of the 2^n intervals of K_n contains a point of K_C.

Let $d < 1$ and choose an $n \in \mathbb{N}$ so that

$$3^{-n} \leq d < 3^{-(n-1)}.$$

We find an upper bound for $N_d(K_C)$ by observing that each of the 2^n line segments making up K_n intersects at most six boxes from G_d (Figure 3.8), and so

$$N_d(K_C) \leq N_d(K_n) \leq 6 \times 2^n = 3 \times 2^{n+1} \quad \text{since } K_C \subseteq K_n.$$

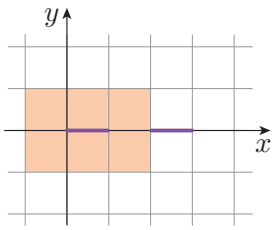

Figure 3.8

To find a lower bound for $N_d(K_C)$, we note that each interval in K_n contains a point of K_C, and each box in G_d can intersect at most two intervals in K_n (Figure 3.9), so

$$N_d(K_C) \geq \frac{2^n}{2} = 2^{n-1}.$$

Hence

$$2^{n-1} \leq N_d(K_C) \leq 3 \times 2^{n+1},$$

and so

$$(n-1)\log 2 \leq \log N_d(K_C) \leq (n+1)\log 2 + \log 3. \qquad (2)$$

Now, since $3^{-n} \leq d < 3^{-(n-1)}$,

$$-n\log 3 \leq \log d < -(n-1)\log 3,$$

and so, for $n \geq 2$,

$$\frac{1}{n\log 3} \leq \frac{1}{-\log d} < \frac{1}{(n-1)\log 3}.$$

Thus for $n \geq 2$ and $3^{-n} \leq d < 3^{-(n-1)}$, we can combine these inequalities with (2) to obtain

$$\frac{(n-1)\log 2}{n\log 3} \leq \frac{\log N_d(K_C)}{-\log d} < \frac{(n+1)\log 2 + \log 3}{(n-1)\log 3}.$$

Now

$$\lim_{n\to\infty} \frac{(n-1)\log 2}{n\log 3} = \lim_{n\to\infty} \left(\frac{1-(1/n)}{1}\right)\frac{\log 2}{\log 3} = \frac{\log 2}{\log 3}$$

and

$$\lim_{n\to\infty} \frac{(n+1)\log 2 + \log 3}{(n-1)\log 3} = \lim_{n\to\infty} \left(\frac{1+(1/n)}{1-(1/n)}\right)\frac{\log 2}{\log 3} + \lim_{n\to\infty}\frac{1}{n-1} = \frac{\log 2}{\log 3}.$$

Since $n \to \infty$ as $d \to 0$, we deduce that

$$\dim K_C = \lim_{d\to 0} \frac{\log N_d(K_C)}{-\log d} = \frac{\log 2}{\log 3}.$$

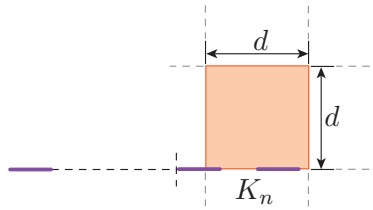

Figure 3.9

Exercise 3.2

Let $S\colon \mathcal{K}(\mathbb{R}) \to \mathcal{K}(\mathbb{R})$ be defined by

$$S(K) = S_1(K) \cup S_2(K),$$

where $S_1(x) = \frac{1}{4}(x+1)$ and $S_2(x) = \frac{1}{4}(x+3)$.

Let $K_0 = [0,1]$ and $K_n = S^n(K_0)$, for each $n \in \mathbb{N}$.

Determine $\dim K_S$, where $K_S = \bigcap_{n\in\mathbb{N}} K_n$.

3.2 The open set condition

We have now found a way of defining the dimension of an intricate set. As you have seen, there may be many calculations involved in determining the dimension of a set. We now wish to calculate the dimensions of the fractals that we constructed in Section 2; each of these is the invariant set of a family of similarities $S_1, S_2, \ldots, S_k$. In this case, there is a simple formula for the dimension of the set, provided that the similarities satisfy a special condition.

Definition 3.4 *Open set condition*

The similarities $S_1, S_2, \ldots, S_k$ satisfy the **open set condition** if there exists a non-empty bounded open set U such that

$$S_i(U) \cap S_j(U) = \emptyset, \quad \text{for } i \neq j, \text{ and } \bigcup_{i=1}^{k} S_i(U) \subseteq U.$$

Remark

The open set condition is a 'disjointness' condition: it says that there is an open set U such that the images of U under distinct similarities are disjoint.

Cantor set

For example, consider the middle-third Cantor set, which is the invariant set of the similarities

$$S_1(x) = \tfrac{1}{3}x \quad \text{and} \quad S_2(x) = \tfrac{1}{3}(x+2).$$

These two similarities satisfy the open set condition, as you can see by taking U to be the open interval $(0,1)$: then $S_1(U) = (0, \frac{1}{3})$ and $S_2(U) = (\frac{2}{3}, 1)$ so

$$S_1(U) \cap S_2(U) = \emptyset$$

and

$$S_1(U) \cup S_2(U) = (0, \tfrac{1}{3}) \cup (\tfrac{2}{3}, 1) \subseteq U.$$

Exercise 3.3

The invariant set of S_1, S_2, S_3 here is the Sierpiński gasket.

Let $S_1, S_2, S_3 \colon \mathbb{R}^2 \to \mathbb{R}^2$ be the similarities defined by

$$S_1(x,y) = \left(\tfrac{1}{2}x, \tfrac{1}{2}y\right),$$
$$S_2(x,y) = \left(\tfrac{1}{2}x + \tfrac{1}{2}, \tfrac{1}{2}y\right),$$
$$S_3(x,y) = \left(\tfrac{1}{2}x + \tfrac{1}{4}, \tfrac{1}{2}y + \tfrac{\sqrt{3}}{4}\right).$$

Show that S_1, S_2 and S_3 satisfy the open set condition.

We now give the main result of this section.

Theorem 3.5

Let K_S be the invariant set of the similarities $S_1, S_2, \ldots, S_k$, and let $c_i \in (0,1)$ be the similarity ratio of S_i, for $1 \leq i \leq k$. If $S_1, S_2, \ldots, S_k$ satisfy the open set condition, then

$$\dim K_S = s,$$

where s is the solution of the equation $\sum_{i=1}^{k} c_i^s = 1$.

Proof (sketch) The proof of this result is fairly complicated, and so we just outline some of the ideas that are involved.

This proof is not assessed.

In particular, we indicate why the box dimension of K_S, if it exists, must be the real number s that satisfies

$$c_1^s + c_2^s + \cdots + c_k^s = 1.$$

To reduce writing, we suppose that we have only two similarities $S_1, S_2 \colon \mathbb{R}^2 \to \mathbb{R}^2$, with similarity ratios c_1 and c_2, respectively.

Since K_S is the invariant set of the similarities, we know that

$$K_S = S_1(K_S) \cup S_2(K_S).$$

The key observation underlying the proof is that each $S_i(K_S)$ is just a copy of the set K_S scaled by a factor c_i, and so there should be a relationship between the number of boxes of side $r > 0$ that intersect K_S, and the number of boxes of side $c_i r$ that intersect $S_i(K_S)$.

If we look at a grid G_r and its image under a similarity S_i, then we obtain a grid of boxes of side length $c_i r$. Hence, if $N_r(K_S)$ boxes from G_r meet K_S, then the same number of boxes from the transformed grid meet $S_i(K_S)$. We would like to conclude that $N_{c_i r}(S_i(K_S)) = N_r(K_S)$, but, unfortunately, the image of a box from G_r may not be a box from the grid $G_{c_i r}$ (as indicated in Figure 3.10), and so the equality need not hold.

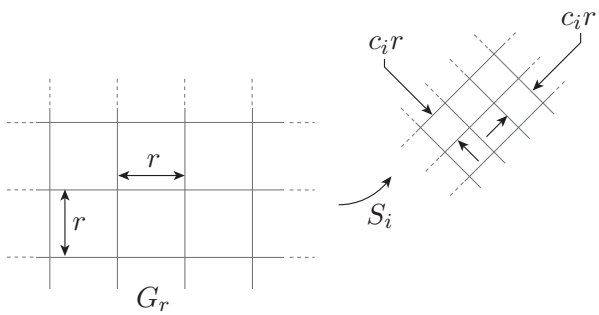

Figure 3.10

However, it is possible to show that there are real numbers $0 < a \leq 1 \leq b$ so that, for any non-empty compact set K and similarity S with similarity ratio c,

$$aN_r(K) \leq N_{cr}(S(K)) \leq bN_r(K). \tag{3}$$

Now

$$K_S = S_1(K_S) \cup S_2(K_S),$$

and if we assume that this union is disjoint, then there is a $\delta > 0$ such that for all $\mathbf{x} \in S_1(K_S)$ and $\mathbf{y} \in S_2(K_S)$,

$$d^{(2)}(\mathbf{x}, \mathbf{y}) \geq \delta.$$

Hence if $d > 0$ is small enough, then a box in $G_d(K_S)$ can only intersect at most one of the sets $S_1(K_S)$ and $S_2(K_S)$, and so

In the full proof, the open set condition plays the role of disjointness and ensures that an equation similar to (4) is still true.

$$N_d(K_S) = N_d(S_1(K_S)) + N_d(S_2(K_S)). \tag{4}$$

Thus, by using (3) with $K = K_S$, $S = S_i$ and $r = d/c_i$ for $i = 1, 2$, we find

$$a(N_{d/c_1}(K_S) + N_{d/c_2}(K_S)) \leq N_d(K_S) \leq b(N_{d/c_1}(K_S) + N_{d/c_2}(K_S)).$$

If we now suppose that $s = \dim K_S$ exists, and in particular that there is a positive number m such that

$$N_d(K_S) = m\left(\frac{1}{d}\right)^s \quad \text{for } 0 < d \leq 1,$$

then we deduce that, for sufficiently small $d > 0$,

$$amd^{-s}\left(c_1^s + c_2^s\right) \leq md^{-s} \leq bmd^{-s}\left(c_1^s + c_2^s\right).$$

Hence

$$a\left(c_1^s + c_2^s\right) \leq 1 \leq b\left(c_1^s + c_2^s\right).$$

The composition of two similarities is a similarity whose similarity ratio is the product of the original ratios.

Now consider the similarities $S_i \circ S_j$ for $i, j = 1, 2$. The similarity $S_i \circ S_j$ has similarity ratio $c_i c_j$, and

$$K_S = \bigcup_{i,j=1}^{2} (S_i \circ S_j)(K_S).$$

Hence, by repeating the preceding argument for the four similarities $S_1 \circ S_1,\ S_1 \circ S_2,\ S_2 \circ S_1,\ S_2 \circ S_2$, we find

$$\begin{aligned} a\big((c_1c_1)^s + (c_1c_2)^s + (c_2c_1)^s + (c_2c_2)^s\big) &\leq 1 \\ &\leq b\big((c_1c_1)^s + (c_1c_2)^s + (c_2c_1)^s + (c_2c_2)^s\big). \end{aligned}$$

But

$$(c_1c_1)^s + (c_1c_2)^s + (c_2c_1)^s + (c_2c_2)^s = (c_1^s + c_2^s)^2,$$

and so

$$a\left(c_1^s + c_2^s\right)^2 \leq 1 \leq b\left(c_1^s + c_2^s\right)^2.$$

Similarly, for each $n \in \mathbb{N}$, by considering the similarities $S_{i_1} \circ S_{i_2} \circ \cdots \circ S_{i_n}$ for $1 \leq i_1, i_2, \ldots, i_n \leq 2$, we find that

$$a\left(c_1^s + c_2^s\right)^n \leq 1 \leq b\left(c_1^s + c_2^s\right)^n.$$

Thus

$$\tfrac{1}{b} \leq (c_1^s + c_2^s)^n \leq \tfrac{1}{a} \quad \text{for each } n \in \mathbb{N}. \tag{5}$$

If $c_1^s + c_2^s > 1$ then $(c_1^s + c_2^s)^n > 1/a$ for large enough n, and if $c_1^s + c_2^s < 1$ then $(c_1^s + c_2^s)^n < 1/b$ for large enough n. Thus to satisfy (5), we must have $c_1^s + c_2^s = 1$, as claimed.

The complete argument to show that the box dimension exists is based on similar ideas but requires more work; we omit the details. ■

Worked Exercise 3.6

Let K_0 and K_1 be as shown in Figure 3.11.

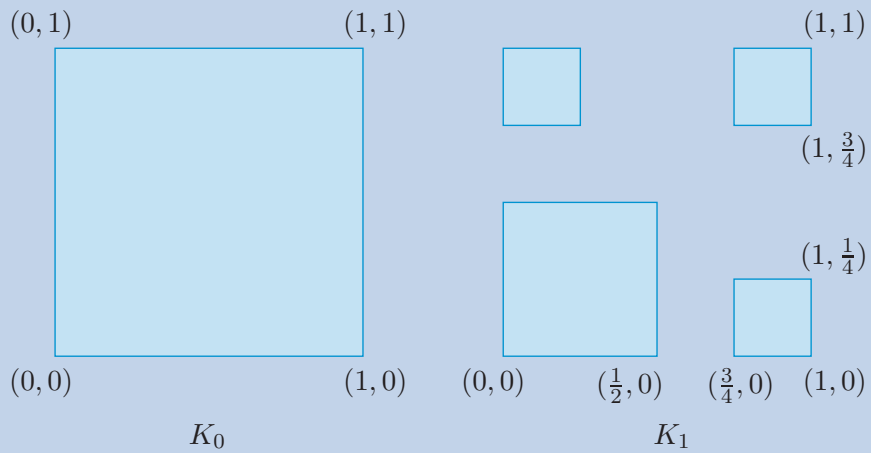

Figure 3.11

(a) Determine similarities S_1, S_2, S_3, S_4 such that $K_1 = \bigcup_{i=1}^4 S_i(K_0)$.

(b) Let $K_{n+1} = \bigcup_{i=1}^4 S_i(K_n)$, for each $n \in \mathbb{N}$. Sketch the set K_2.

(c) Show that $K = \lim_{n\to\infty} K_n$ exists and is the invariant set of S_1, S_2, S_3, S_4.

(d) Show that S_1, S_2, S_3, S_4 satisfy the open set condition.

(e) Determine $\dim K$.

Solution

(a) A natural choice of similarities is:

$$S_1(x, y) = \left(\tfrac{1}{2}x, \tfrac{1}{2}y\right), \qquad S_2(x, y) = \left(\tfrac{1}{4}x + \tfrac{3}{4}, \tfrac{1}{4}y\right),$$
$$S_3(x, y) = \left(\tfrac{1}{4}x + \tfrac{3}{4}, \tfrac{1}{4}y + \tfrac{3}{4}\right), \qquad S_4(x, y) = \left(\tfrac{1}{4}x, \tfrac{1}{4}y + \tfrac{3}{4}\right).$$

(b) See Figure 3.12.

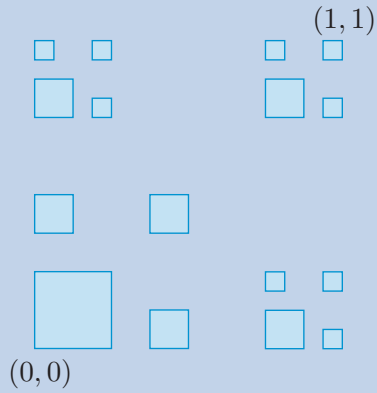

Figure 3.12

(c) Let $S\colon \mathcal{K}(\mathbb{R}^2) \to \mathcal{K}(\mathbb{R}^2)$ be defined by

$$S(K) = \bigcup_{i=1}^{4} S_i(K).$$

Then $K_{n+1} = S(K_n)$, and hence $K_n = S^n(K_0)$ for each $n \in \mathbb{N}$. Since K_0 is non-empty and compact, it follows from Theorem 2.15 that $K = \lim_{n\to\infty} S^n(K_0)$ exists and is the invariant set of S_1, S_2, S_3, S_4.

(d) The following diagram shows that S_1, S_2, S_3, S_4 satisfy the open set condition.

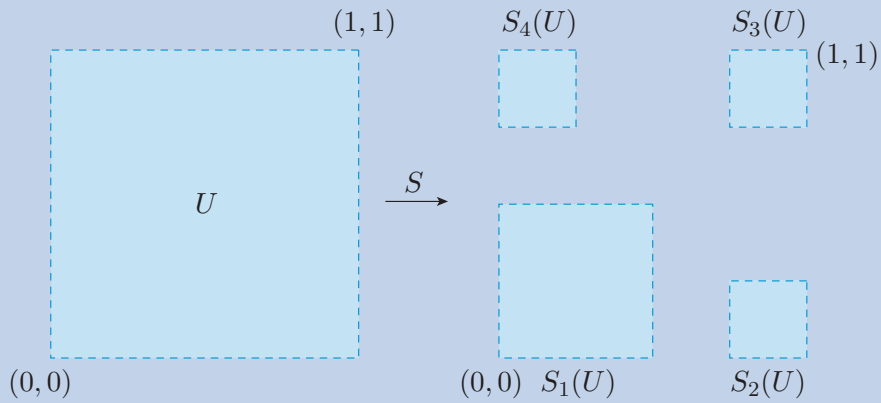

Figure 3.13

(e) It follows from parts (c) and (d) that the conditions of Theorem 3.5 are satisfied. Since S_1 has similarity ratio $\frac{1}{2}$ and each of the similarities S_2, S_3 and S_4 has similarity ratio $\frac{1}{4}$, the box dimension of K is the number s for which

$$\left(\tfrac{1}{2}\right)^s + 3\left(\tfrac{1}{4}\right)^s = 1.$$

By substituting $x = \left(\frac{1}{2}\right)^s$, we obtain

$$x + 3x^2 = 1,$$

which has solutions $x = \frac{1}{6}(-1 + \sqrt{13})$ and $x = \frac{1}{6}(-1 - \sqrt{13})$.

But x must be positive, and so

$$\left(\tfrac{1}{2}\right)^s = \tfrac{1}{6}(-1 + \sqrt{13}).$$

Thus

$$\dim K = s = \frac{\log \frac{1}{6}(-1 + \sqrt{13})}{-\log 2} = \frac{\log 6 - \log(-1 + \sqrt{13})}{\log 2} \simeq 1.203\ldots.$$

The formula for the dimension of a set satisfying the hypotheses of Theorem 3.5 is particularly simple when all the similarities involved have the same similarity ratio – this is the case for each of the fractals in Section 1. For suppose that K_S is the invariant set of the similarities $S_1, S_2, \ldots, S_k$, each with similarity ratio $c \in (0, 1)$, and that $S_1, S_2, \ldots, S_k$ satisfy the open set condition. It follows from Theorem 3.5 that $\dim K_S = s$, where s satisfies

$$\sum_{i=1}^{k} c^s = kc^s = 1,$$

so

$$k = c^{-s}.$$

Taking logarithms, we obtain

$$\log k = -s \log c,$$

that is,

$$s = \frac{\log k}{-\log c}.$$

We have proved the following result.

Theorem 3.7

Let K_S be the invariant set of the similarities $S_1, S_2, \ldots, S_k$, each with similarity ratio $c \in (0, 1)$. If $S_1, S_2, \ldots, S_k$ satisfy the open set condition, then

$$\dim K_S = \frac{\log k}{-\log c}.$$

This theorem gives a simple way of calculating the box dimension of a large class of fractals.

Sierpiński gasket

See Exercise 1.7.

The Sierpiński gasket K_G is the invariant set of the similarities

$$S_1(x, y) = \left(\tfrac{1}{2}x, \tfrac{1}{2}y\right), \qquad S_2(x, y) = \left(\tfrac{1}{2}x + \tfrac{1}{2}, \tfrac{1}{2}y\right),$$

$$S_3(x, y) = \left(\tfrac{1}{2}x + \tfrac{1}{4}, \tfrac{1}{2}y + \tfrac{\sqrt{3}}{4}\right).$$

In Exercise 3.3 you showed that S_1, S_2 and S_3 satisfy the open set condition. Since each of these similarities has similarity ratio $\frac{1}{2}$, it follows from Theorem 3.7 that

$$\dim K_G = \frac{\log 3}{-\log \frac{1}{2}} = \frac{\log 3}{\log 2} \simeq 1.584\ldots.$$

So the dimension of the Sierpiński gasket lies between 1 and 2, as we predicted.

von Koch curve

The von Koch curve K_V is the invariant set of four similarities.

S_1 scales by a factor of $\frac{1}{3}$.

S_2 scales by a factor of $\frac{1}{3}$, rotates anticlockwise by $\frac{\pi}{3}$ about $(0, 0)$ and translates by $\left(\frac{1}{3}, 0\right)$.

S_3 scales by a factor of $\frac{1}{3}$, rotates clockwise by $\frac{\pi}{3}$ about the point $(0, 0)$, and translates by $(\frac{1}{2}, \frac{\sqrt{3}}{6})$.

S_4 scales by a factor of $\frac{1}{3}$ and translates by $\left(\frac{2}{3}, 0\right)$.

Here, U is the interior of the triangle based on $[0, 1]$ with height $\frac{\sqrt{3}}{6}$.

Figure 3.14 shows that these similarities satisfy the open set condition.

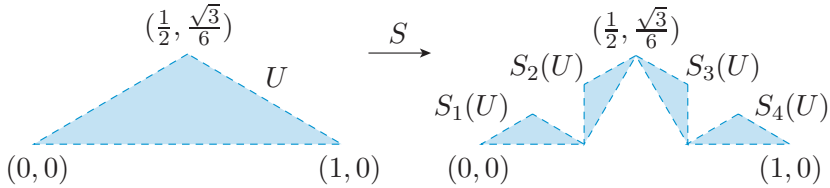

Figure 3.14

Since each of these four similarities has similarity ratio $\frac{1}{3}$, it follows from Theorem 3.7 that

$$\dim K_V = \frac{\log 4}{-\log \frac{1}{3}} = \frac{\log 4}{\log 3} \simeq 1.261\ldots.$$

Thus the dimension of the von Koch curve lies between 1 and 2, as we expected.

So, the three examples of fractals that you met in Section 1 – the middle-third Cantor set, the Sierpiński gasket and the von Koch curve – all satisfy the conditions of Theorem 3.7 and their dimensions can be calculated easily.

We now look at some other examples of fractals for which this is true.

Worked Exercise 3.8

Let K_0 and K_1 be the curves shown in Figure 3.15.

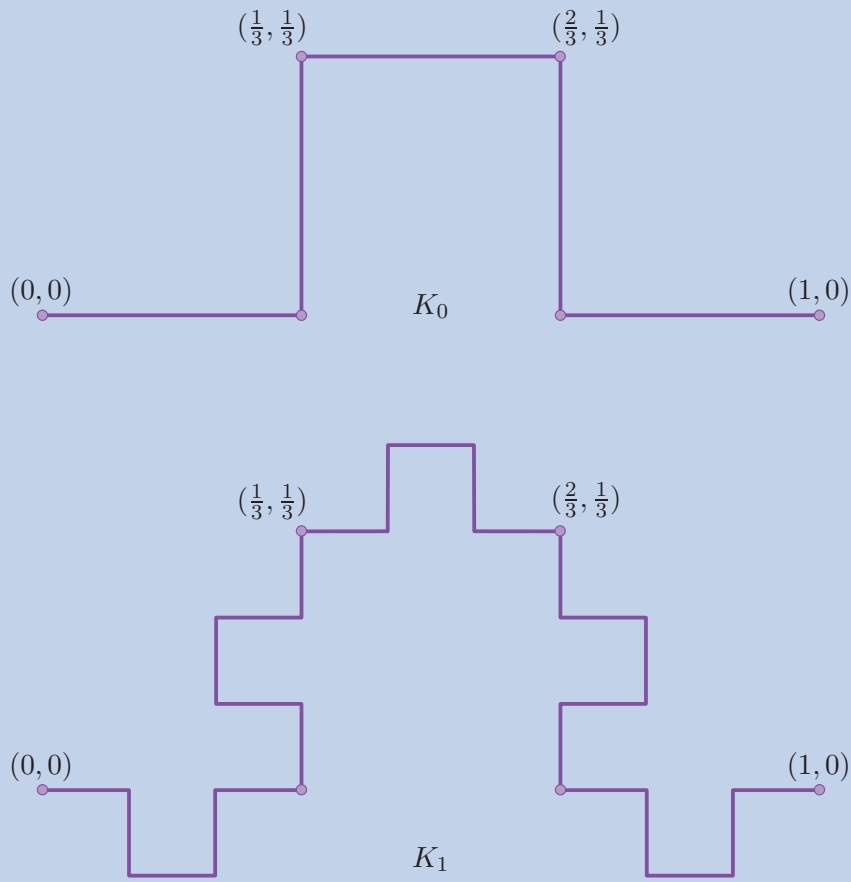

Figure 3.15

(a) Determine similarities $S_1, S_2, \ldots, S_5$ such that $K_1 = \bigcup_{i=1}^{5} S_i(K_0)$.

(b) Let $K_{n+1} = \bigcup_{i=1}^{5} S_i(K_n)$, for each $n \in \mathbb{N}$. Sketch the set K_2.

(c) Show that $K = \lim_{n\to\infty} K_n$ exists and is the invariant set of $S_1, S_2, \ldots, S_5$.

(d) Show that $S_1, S_2, \ldots, S_5$ satisfy the open set condition.

(e) Determine $\dim K$.

Solution

(a) A natural choice of similarities is

$$S_1(x,y) = \left(\tfrac{1}{3}x, -\tfrac{1}{3}y\right), \qquad S_2(x,y) = \left(-\tfrac{1}{3}y + \tfrac{1}{3}, \tfrac{1}{3}x\right),$$
$$S_3(x,y) = \left(\tfrac{1}{3}x + \tfrac{1}{3}, \tfrac{1}{3}y + \tfrac{1}{3}\right), \qquad S_4(x,y) = \left(\tfrac{1}{3}y + \tfrac{2}{3}, -\tfrac{1}{3}x + \tfrac{1}{3}\right),$$
$$S_5(x,y) = \left(\tfrac{1}{3}x + \tfrac{2}{3}, -\tfrac{1}{3}y\right).$$

(b) See Figure 3.16.

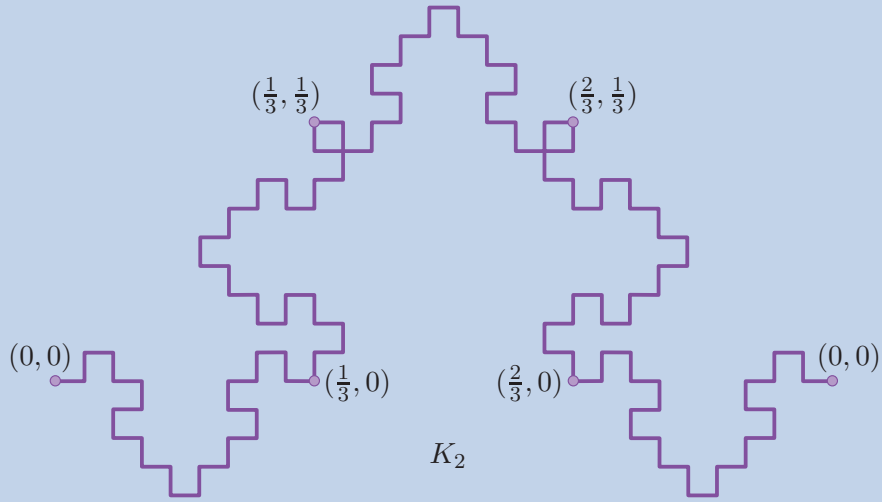

Figure 3.16

(c) Let $S\colon \mathcal{K}(\mathbb{R}^2) \to \mathcal{K}(\mathbb{R}^2)$ be defined by

$$S(K) = \bigcup_{i=1}^{5} S_i(K).$$

Then $K_{n+1} = S(K_n)$, and hence $K_n = S^n(K_0)$ for each $n \in \mathbb{N}$. Since K_0 is non-empty and compact, it follows from Theorem 2.15 that $K = \lim_{n\to\infty} S^n(K_0)$ exists and is the invariant set of $S_1, S_2, \ldots, S_5$.

(d) Figure 3.17 shows that $S_1, S_2, \ldots, S_5$ satisfy the open set condition.

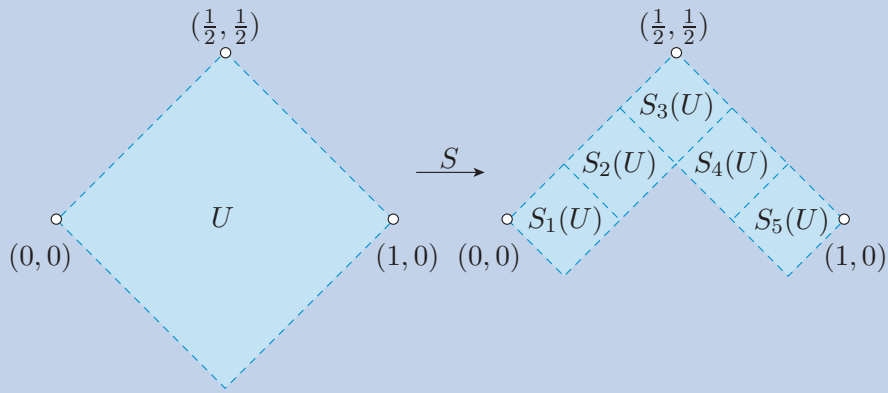

Figure 3.17

(e) It follows from parts (c) and (d) that the conditions of Theorem 3.7 are satisfied. Since each of the similarities $S_1, S_2, \ldots, S_5$ has similarity ratio $\frac{1}{3}$, we have

$$\dim K = \frac{\log 5}{-\log \frac{1}{3}} = \frac{\log 5}{\log 3} \simeq 1.464\ldots.$$

Exercise 3.4

Let K_0 and K_1 be the sets shown in Figure 3.18.

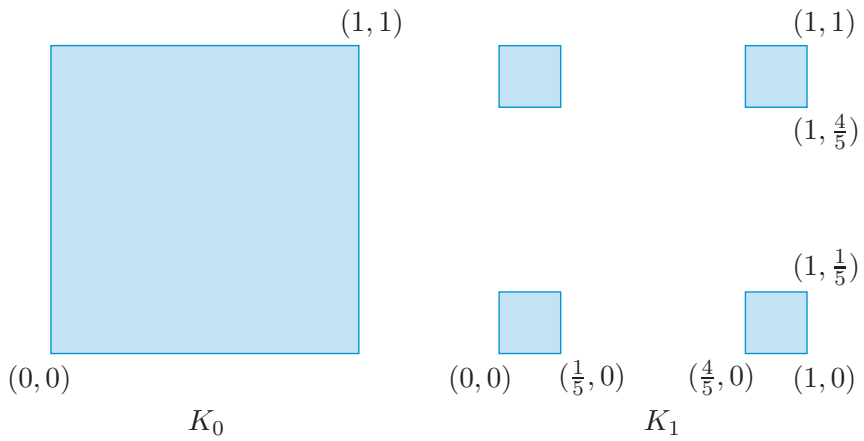

Figure 3.18

(a) Determine similarities S_1, S_2, S_3, S_4 such that $K_1 = \bigcup_{i=1}^{4} S_i(K_0)$.

(b) Let $K_{n+1} = \bigcup_{i=1}^{4} S_i(K_n)$, for each $n \in \mathbb{N}$. Sketch the set K_2.

(c) Show that $K = \lim_{n\to\infty} K_n$ exists and is the invariant set of S_1, S_2, S_3, S_4.

(d) Show that S_1, S_2, S_3, S_4 satisfy the open set condition.

(e) Determine $\dim K$.

Conclusion

In this chapter, we have seen how general results concerning metric spaces can be used to great effect in the study of fractals. In particular, given a family of similarities $S_1, S_2, \ldots, S_k$, we have seen how the Contraction Mapping Theorem enables us to construct a non-empty compact set that is invariant for these similarities. Many sets constructed in this way are fractals, and a good estimate of their size can be obtained by using the formula given in Theorem 3.5.

This is also the last chapter of Book F, and we hope that it has given you a flavour of how ideas and results from the theory of metric spaces can be applied to study other areas of mathematics.

We have travelled a long way from our initial study of continuous functions on the real line, and have gained a deeper understanding of the role that

metric spaces play in modern mathematics. The theory of metric spaces is a vital part of every pure mathematician's toolkit.

Learning outcomes

We list those topics on which we may set assessment questions to test your understanding of the chapter.

After working through this chapter you should be able to:

(a) Describe standard examples of fractal sets, such as the Cantor set, Sierpiński gasket and the von Koch curve.

(b) Calculate the Hausdorff distance between given non-empty compact sets.

(c) Appreciate how the Contraction Mapping Theorem can be used to define fractal sets as the invariant sets of a finite collection of similarities.

(d) Determine the box dimension of invariant sets that satisfy the open set condition.

Solutions to exercises

Solution to Exercise 1.1

The set K_n consists of 2^n line segments, each of length $\frac{1}{3^n}$. Their total length is $L_n = \frac{2^n}{3^n} = \left(\frac{2}{3}\right)^n$, and $\lim_{n\to\infty} L_n = 0$.

Solution to Exercise 1.2

(a) S has similarity ratio $\frac{1}{2}$.

(b) S has similarity ratio $\frac{1}{4}$.

Solution to Exercise 1.3

There are many possibilities; for example, the similarities

$$S_1(x) = \tfrac{1}{2}x \quad \text{and} \quad S_2(x) = \tfrac{1}{2}(x+1)$$

have similarity ratio $\frac{1}{2}$ and

$$S_1([0,1]) \cup S_2([0,1]) = \left[0, \tfrac{1}{2}\right] \cup \left[\tfrac{1}{2}, 1\right] = [0,1].$$

Solution to Exercise 1.4

(a) The set K_2 is shown below.

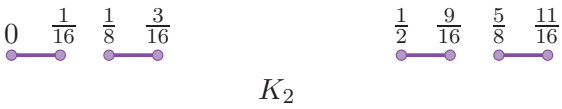

K_2

(b) The set K_n consists of 2^n line segments, each of length $\left(\frac{1}{4}\right)^n$. So the length of K_n is $L_n = \frac{2^n}{4^n} = \left(\frac{1}{2}\right)^n$, and $\lim_{n\to\infty} L_n = 0$.

(c) A natural choice of similarities is

$$S_1(x) = \tfrac{1}{4}x \quad \text{and} \quad S_2(x) = \tfrac{1}{4}x + \tfrac{1}{2},$$

and each has similarity ratio $\frac{1}{4}$.

Solution to Exercise 1.5

Each of the 3^n triangles that form K_n has sides of length $\left(\frac{1}{2}\right)^n$, so each of these triangles has boundary length $3\left(\frac{1}{2}\right)^n$. It follows that the total length of their boundaries is $3^n \times \frac{3}{2^n} = 3\left(\frac{3}{2}\right)^n$.

The boundary of each of the finitely many triangles in K_n belongs to the Sierpiński gasket, and the total length of their boundaries is $3\left(\frac{3}{2}\right)^n$. Since $\left(3\left(\frac{3}{2}\right)^n\right)$ is a divergent sequence, then we may choose $n \in \mathbb{N}$ so that the total length of the boundaries of the triangles in K_n exceeds any given $N \in \mathbb{R}$.

Solution to Exercise 1.6

The set K_n consists of 3^n triangles, each of area

$$\frac{1}{2} \times \frac{\sqrt{3}/2}{2^n} \times \frac{1}{2^n} = \frac{\sqrt{3}}{4^{n+1}}.$$

So the area of K_n is $\frac{3^n\sqrt{3}}{4^{n+1}} = \frac{\sqrt{3}}{4}\left(\frac{3}{4}\right)^n$. Since $K_{\mathrm{G}} \subset K_n$, the area of K_{G} is at most $\frac{\sqrt{3}}{4}\left(\frac{3}{4}\right)^n$ for each $n \in \mathbb{N}$. Since $\lim_{n\to\infty} \frac{\sqrt{3}}{4}\left(\frac{3}{4}\right)^n = 0$, the Sierpiński gasket has zero area.

Solution to Exercise 1.7

A natural choice of similarities is

$$S_1(x, y) = \left(\tfrac{1}{2}x, \tfrac{1}{2}y\right), \qquad S_2(x, y) = \left(\tfrac{1}{2}x + \tfrac{1}{2}, \tfrac{1}{2}y\right),$$

$$S_3(x, y) = \left(\tfrac{1}{2}x + \tfrac{1}{4}, \tfrac{1}{2}y + \tfrac{\sqrt{3}}{4}\right),$$

and they each have similarity ratio $\frac{1}{2}$.

Solution to Exercise 1.8

(a) The set K_2 is shown below.

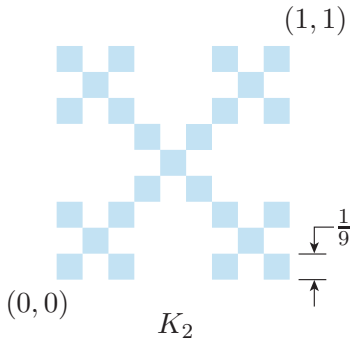

(b) Each set K_n is a finite union of closed bounded squares, and so is non-empty and compact. Since $K_{n+1} \subseteq K_n$, the set K is the intersection of a nested sequence of non-empty compact sets, and so is non-empty and compact.

The set K_n consists of 5^n closed squares, each with sides of length $\left(\frac{1}{3}\right)^n$ and area $\left(\frac{1}{9}\right)^n$, so the total area of K_n is $\left(\frac{5}{9}\right)^n$. Since $K \subset K_n$, the area of K is at most $\left(\frac{5}{9}\right)^n$ for each $n \in \mathbb{N}$. Since $\lim_{n\to\infty}\left(\frac{5}{9}\right)^n = 0$, K has zero area.

(c) A natural choice of similarities is

$$S_1(x, y) = \left(\tfrac{1}{3}x, \tfrac{1}{3}y\right), \qquad S_2(x, y) = \left(\tfrac{1}{3}x + \tfrac{2}{3}, \tfrac{1}{3}y\right),$$

$$S_3(x, y) = \left(\tfrac{1}{3}x, \tfrac{1}{3}y + \tfrac{2}{3}\right), \qquad S_4(x, y) = \left(\tfrac{1}{3}x + \tfrac{2}{3}, \tfrac{1}{3}y + \tfrac{2}{3}\right),$$

$$S_5(x, y) = \left(\tfrac{1}{3}x + \tfrac{1}{3}, \tfrac{1}{3}y + \tfrac{1}{3}\right),$$

and each has similarity ratio $\frac{1}{3}$.

Solution to Exercise 1.9

The set K_1 consists of four line segments of length $\frac{1}{3}$. When constructing K_n, each line segment in K_{n-1} is replaced by four line segments each of $\frac{1}{3}$ of the length. So K_n is a connected curve consisting of 4^n line segments each of length $\left(\frac{1}{3}\right)^n$. Using the method given in the question, we conclude that K_{V} has length at least $\left(\frac{4}{3}\right)^n$, for each $n \in \mathbb{N}$. Since $\left(\left(\frac{4}{3}\right)^n\right)$ is a divergent sequence, K_{V} does not have finite length.

Solution to Exercise 2.1

(a) $(-1, 1]$ is not closed and so is not compact – so it does not belong to $\mathcal{K}(\mathbb{R})$.

(b) $[0, 2] \in \mathcal{K}(\mathbb{R})$.

(c) $[0, \infty)$ is not bounded and so is not compact – so it does not belong to $\mathcal{K}(\mathbb{R})$.

(d) $\{1\} \in \mathcal{K}(\mathbb{R})$.

(e) $[0, 2] \cap [4, 6] = \emptyset \notin \mathcal{K}(\mathbb{R})$.

Solution to Exercise 2.2

(a) $d^{(1)}(1, [2, 3]) = d^{(1)}(1, 2) = 1$.

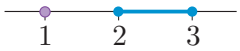

(b) $d^{(1)}(1, [0, 2]) = d^{(1)}(1, 1) = 0$.

(c) $d^{(2)}((0, 3), \{(x, y) : x^2 + y^2 \leq 1\}) = d^{(2)}((0, 3), (0, 1)) = 2$.

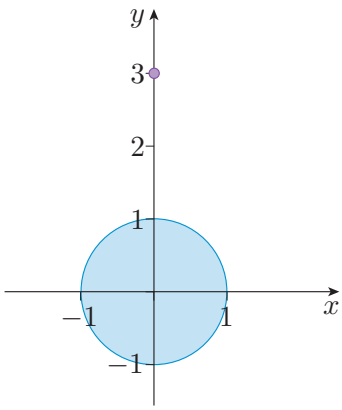

(d) $d^{(2)}((2,1),\{(3,y):|y|\leq 1\}) = d^{(2)}((2,1),(3,1)) = 1.$

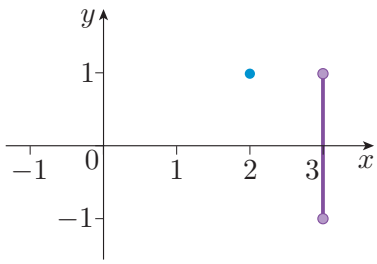

Solution to Exercise 2.3

Let $J \subseteq K$. Then

$$\begin{aligned} d^{(m)}(\mathbf{x},K) &= \min\{d^{(m)}(\mathbf{x},\mathbf{k}) : \mathbf{k}\in K\} \\ &\leq \min\{d^{(m)}(\mathbf{x},\mathbf{k}) : \mathbf{k}\in J\cap K\} \\ &= \min\{d^{(m)}(\mathbf{x},\mathbf{k}) : \mathbf{k}\in J\} = d^{(m)}(\mathbf{x},J). \end{aligned}$$

Solution to Exercise 2.4

We begin by sketching the sets J and K.

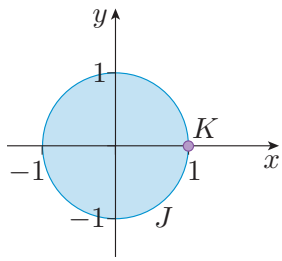

The point in J that is furthest from K is $(-1,0)$, so

$$\begin{aligned} d^{(2)}(J,K) &= d^{(2)}((-1,0),K) \\ &= d^{(2)}((-1,0),(1,0)) = 2. \end{aligned}$$

There is only one point in K, so

$$d^{(2)}(K,J) = d^{(2)}((1,0),J) = d^{(2)}((1,0),(1,0)) = 0.$$

Solution to Exercise 2.5

(a) $$\begin{aligned} d^{(m)}(J,K) &= \max\{d^{(m)}(\mathbf{x},K) : \mathbf{x}\in J\} \\ &\leq \max\{d^{(m)}(\mathbf{x},J) : \mathbf{x}\in J\}, \end{aligned}$$

by Lemma 2.6, since $J \subseteq K$.

By Lemma 2.5, $d^{(m)}(\mathbf{x},J) = 0$ if $\mathbf{x}\in J$, and so

$$d^{(m)}(J,K) = 0.$$

(b) $$\begin{aligned} d^{(m)}(L,K) &= \max\{d^{(m)}(\mathbf{x},K) : \mathbf{x}\in L\} \\ &\leq \max\{d^{(m)}(\mathbf{x},J) : \mathbf{x}\in L\}, \text{ by Lemma 2.6} \\ &= d^{(m)}(L,J). \end{aligned}$$

Solution to Exercise 2.6

The solution to Exercise 2.4 shows that $d^{(2)}(J,K) = 2$ and $d^{(2)}(K,J) = 0$. So

$$d_{\mathrm{H}}(J,K) = \max\{2,0\} = 2.$$

Solution to Exercise 2.7

(a) $$\begin{aligned} S([0,2]) &= S_1([0,2]) \cup S_2([0,2]) \\ &= \left[0,\tfrac{2}{3}\right] \cup \left[\tfrac{2}{3},\tfrac{4}{3}\right] = \left[0,\tfrac{4}{3}\right]. \end{aligned}$$

This set belongs to $\mathcal{K}(\mathbb{R})$.

(b) $$\begin{aligned} S([-2,2]) &= S_1([-2,2]) \cup S_2([-2,2]) \\ &= \left[-\tfrac{2}{3},\tfrac{2}{3}\right] \cup \left[0,\tfrac{4}{3}\right] = \left[-\tfrac{2}{3},\tfrac{4}{3}\right]. \end{aligned}$$

This set belongs to $\mathcal{K}(\mathbb{R})$.

Solution to Exercise 2.8

(a)

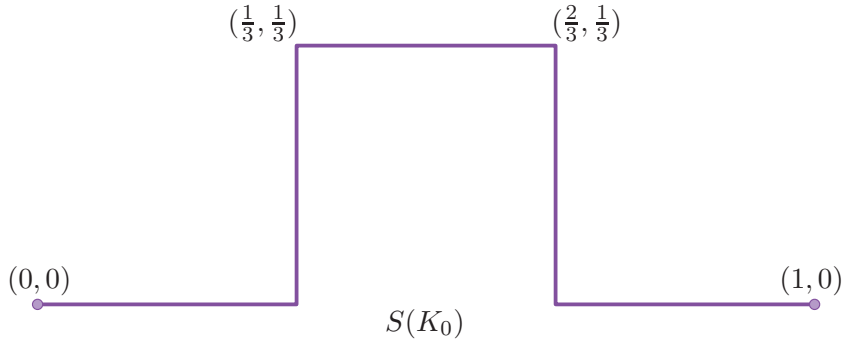

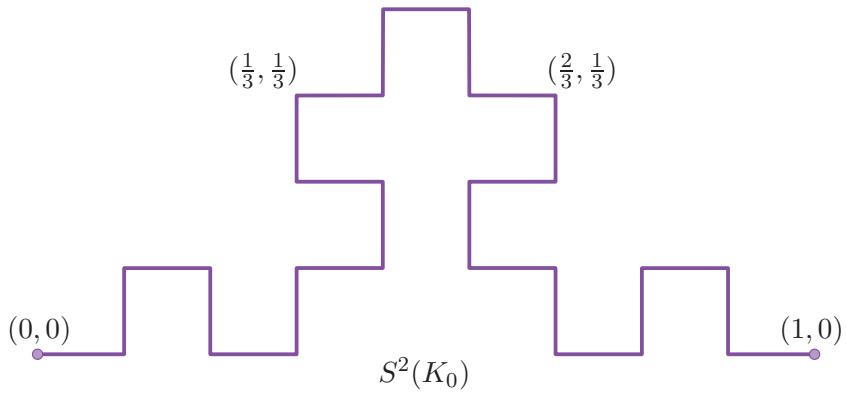

(b) Let $K_n = S^n(K_0)$, for each $n \in \mathbb{N}$.

The set K_1 consists of five line segments each of length $\frac{1}{3}$ and K_2 consists of 25 line segments each of length $(\frac{1}{3})^2$. In constructing K_n, each line segment of K_{n-1} is replaced by five line segments of one-third the length, so K_n consists of 5^n line segments of length $\frac{1}{3^n}$ and so has length $\left(\frac{5}{3}\right)^n$.

(c) Since $K_0 \in \mathcal{K}(\mathbb{R}^2)$, and S is a contraction mapping with contraction ratio $\frac{1}{3}$, it follows from Theorem 2.15 that there is a unique set $K_S \in \mathcal{K}(\mathbb{R}^2)$ such that $K_S = \lim_{n\to\infty} S^n(K_0)$.

(This set satisfies

$$K_S = S(K_S) = \bigcup_{i=1}^{5} S_i(K_S)$$

and so is self-similar.)

Solution to Exercise 2.9

$S_1(\emptyset) = S_2(\emptyset) = \emptyset$, and so $S_1(\emptyset) \cup S_2(\emptyset) = \emptyset$.

$S_1(\mathbb{R}) = S_2(\mathbb{R}) = \mathbb{R}$, and so $S_1(\mathbb{R}) \cup S_2(\mathbb{R}) = \mathbb{R}$.

$S_1([0,\infty)) = [0,\infty)$ and $S_2([0,\infty)) = \left[\frac{2}{3},\infty\right)$ and so $S_1([0,\infty)) \cup S_2([0,\infty)) = [0,\infty)$.

Solution to Exercise 2.10

The set $[0,2]$ is closed and bounded, and hence is in $\mathcal{K}(\mathbb{R})$. Also,

$$S([0,2]) = \left[0,\tfrac{2}{3}\right] \cup \left[\tfrac{2}{3},\tfrac{4}{3}\right] = \left[0,\tfrac{4}{3}\right] \subseteq [0,2].$$

It follows from Theorem 2.15 and Lemma 2.17 that $\bigcap_{n\in\mathbb{N}} S^n([0,2])$ is the unique set in $\mathcal{K}(\mathbb{R})$ that is invariant for S_1 and S_2. Thus $\bigcap_{n\in\mathbb{N}} S^n([0,2]) = K_{\mathrm{C}}$.

Solution to Exercise 2.11

(a)

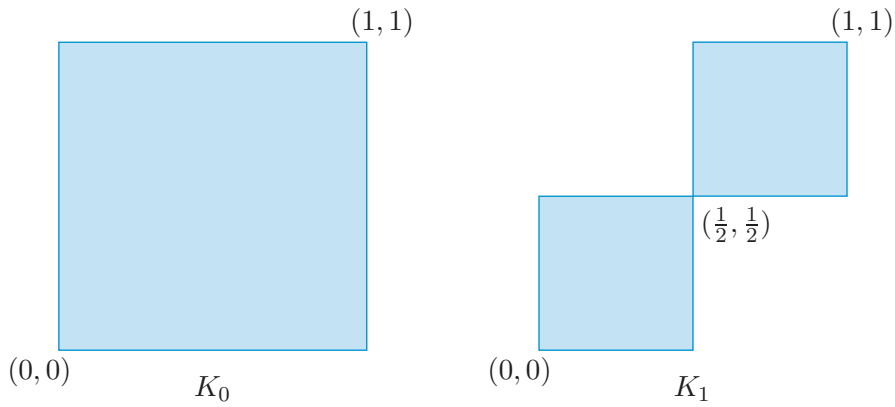

(b) It follows from Theorem 2.15 that $\lim_{n\to\infty} K_n$ is the unique set in $\mathcal{K}(\mathbb{R}^2)$ that is invariant for S. Now if $K' = \{(x,y) : y = x \text{ and } 0 \le x \le 1\}$, then

$$\begin{aligned} S(K') &= S_1(K') \cup S_2(K') \\ &= \{(x,y) : y = x \text{ and } 0 \le x \le \tfrac{1}{2}\} \cup \{(x,y) : y = x \text{ and } \tfrac{1}{2} \le x \le 1\} \\ &= K'. \end{aligned}$$

Thus K' is invariant for S.

Since $K' \in \mathcal{K}(\mathbb{R}^2)$, we deduce that $K' = \lim_{n\to\infty} K_n$.

Solution to Exercise 3.1

(a) $N_{\frac{1}{4}}(A) = 12$.

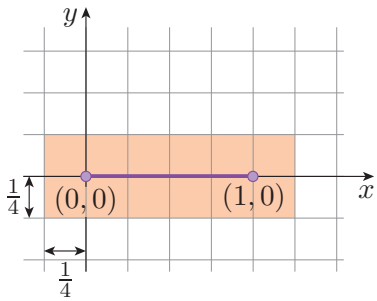

(b) $N_{\frac{1}{4}}(A) = 6$.

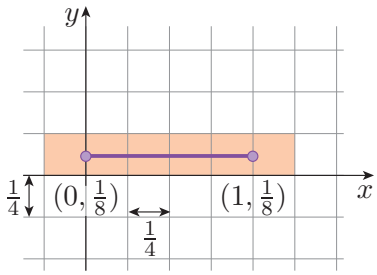

(c) $N_{\frac{1}{4}}(A) = 5$.

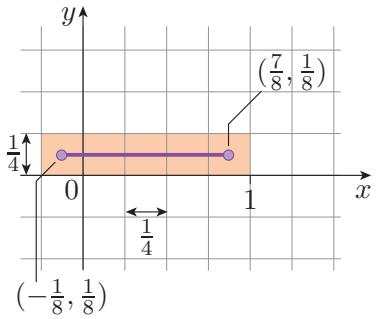

(d) $N_{\frac{1}{4}}(A) = 10$.

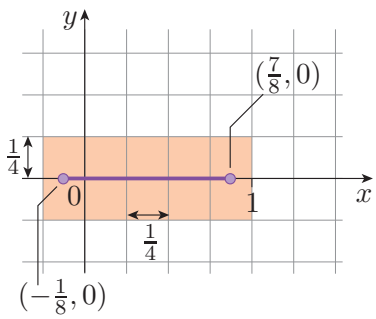

Solution to Exercise 3.2

We follow the technique described in Worked Exercise 3.3.

Our objective is to find

$$\lim_{d\to 0} \frac{\log N_d(K_S)}{-\log d}.$$

For a given $d > 0$, we estimate $N_d(K_S)$ by comparing it with $N_d(K_n)$, where n is chosen so that 4^{-n} is approximately equal to d.

We note first that, at the nth stage of the construction of K_S, the set K_n consists of 2^n line segments each of length $\left(\frac{1}{4}\right)^n$. The set K_n contains K_S, and since the endpoints of the intervals making up K_n are not removed in the construction, each of the 2^n intervals of K_n contains a point of K_S.

Let $d < 1$ and choose an $n \in \mathbb{N}$ so that

$$4^{-n} \leq d < 4^{-(n-1)}.$$

We find an upper bound for $N_d(K_S)$ by observing that each of the 2^n line segments making up K_n intersects at most six boxes from G_d, and so, since $K_S \subseteq K_n$,

$$N_d(K_S) \leq N_d(K_n) \leq 6 \times 2^n = 3 \times 2^{n+1}.$$

To find a lower bound for $N_d(K_S)$, we notice that each interval in K_n contains a point of K_S, and each box in G_d can intersect at most two intervals in K_n, so

$$N_d(K_S) \geq \frac{2^n}{2} = 2^{n-1}.$$

Hence

$$2^{n-1} \leq N_d(K_S) \leq 3 \times 2^{n+1},$$

and so

$$\begin{aligned}(n-1)\log 2 &\leq \log N_d(K_C) \\ &\leq (n+1)\log 2 + \log 3. \end{aligned} \tag{S.1}$$

Now, since $4^{-n} \leq d < 4^{-(n-1)}$,

$$-n \log 4 \leq \log d < -(n-1)\log 4,$$

and so, for $n \geq 2$,

$$\frac{1}{n\log 4} \leq \frac{1}{-\log d} < \frac{1}{(n-1)\log 4}.$$

Thus, for $n \geq 2$ and $4^{-n} \leq d < 4^{-(n-1)}$, we can combine these inequalities with (S.1) to obtain

$$\frac{(n-1)\log 2}{n\log 4} \leq \frac{\log N_d(K_S)}{-\log d} < \frac{(n+1)\log 2 + \log 3}{(n-1)\log 4}.$$

Now

$$\lim_{n\to\infty} \frac{(n-1)\log 2}{n\log 4} = \lim_{n\to\infty} \left(\frac{1-(1/n)}{1}\right) \frac{\log 2}{\log 4}$$
$$= \frac{\log 2}{\log 4} = \frac{1}{2}$$

and

$$\lim_{n\to\infty} \frac{(n+1)\log 2 + \log 3}{(n-1)\log 4}$$
$$= \lim_{n\to\infty} \left(\frac{1+(1/n)}{1-(1/n)}\right) \frac{\log 2}{\log 4} + \lim_{n\to\infty} \left(\frac{1}{n-1}\right) \frac{\log 3}{\log 4}$$
$$= \frac{\log 2}{\log 4} = \frac{1}{2}.$$

Since $n \to \infty$ as $d \to 0$, we deduce

$$\dim K_S = \lim_{d\to 0} \frac{\log N_d(K_{\mathrm{C}})}{-\log d} = \frac{1}{2}.$$

Solution to Exercise 3.3

We take U to be the interior of the equilateral triangle based on the line segment joining $(0,1)$ to $(1,0)$, as shown below.

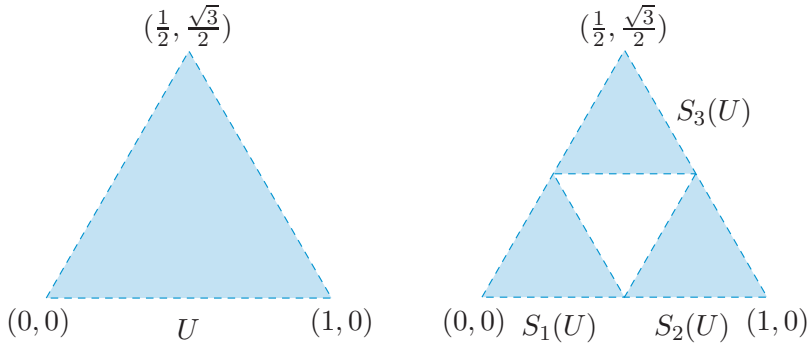

From the figure, we see that the sets $S_1(U)$, $S_2(U)$ and $S_3(U)$ are disjoint and $S_1(U) \cup S_2(U) \cup S_3(U) \subseteq U$, and so the open set condition is satisfied.

Solution to Exercise 3.4

(a) A natural choice of similarities is

$$S_1(x,y) = \left(\tfrac{1}{5}x, \tfrac{1}{5}y\right),$$
$$S_2(x,y) = \left(\tfrac{1}{5}(x+4), \tfrac{1}{5}y\right),$$
$$S_3(x,y) = \left(\tfrac{1}{5}x, \tfrac{1}{5}(y+4)\right),$$
$$S_4(x,y) = \left(\tfrac{1}{5}(x+4), \tfrac{1}{5}(y+4)\right).$$

(b)

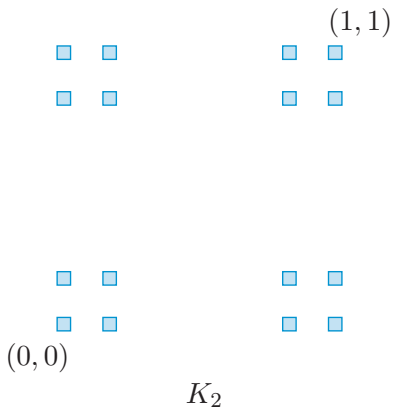

K_2

(c) Let $S\colon \mathcal{K}(\mathbb{R}^2) \to \mathcal{K}(\mathbb{R}^2)$ be defined by

$$S(K) = \bigcup_{i=1}^{4} S_i(K).$$

Then $K_{n+1} = S(K_n)$, and hence $S_n = S^n(K_0)$ for each $n \in \mathbb{N}$. Since K_0 is non-empty and compact, it follows from Theorem 2.15 that $K = \lim_{n\to\infty} S^n(K_0)$ exists and is the invariant set of S_1, S_2, S_3, S_4.

(d) The diagram below shows that S_1, S_2, S_3, S_4 satisfy the open set condition.

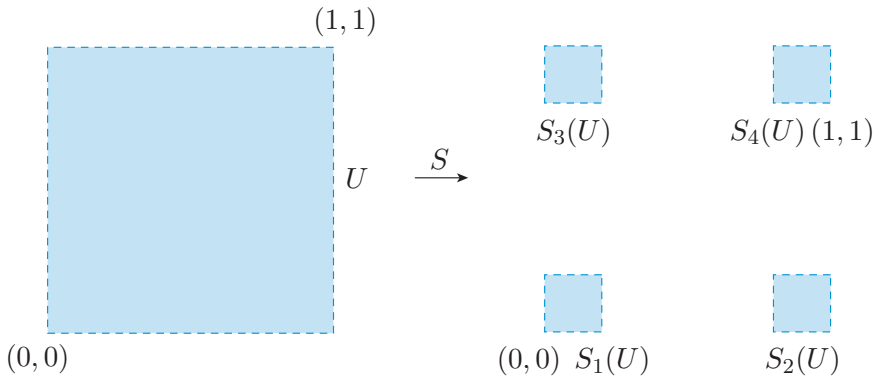

(e) It follows from parts (c) and (d) that the conditions of Theorem 3.7 are satisfied. Since each of the similarities S_1, S_2, S_3, S_4 has similarity ratio $\frac{1}{5}$, we have

$$\dim K = \frac{\log 4}{-\log \frac{1}{5}} = \frac{\log 4}{\log 5}.$$

Index

Arzelá–Ascoli Theorem 93

bounded
  set, ball definition 82
  totally 115
bounded function 71
bounded real sequence 119
bounded set 72, 81
  in a metric space 81
  in a metric space, ball definition 82
Boundedness Theorem 69, 74
  using open covers 107
box dimension 225, 228

$C[0,1]$
  complete 151
  connected 48
  path-connected 47
Cantor sets 194, 219
  construction of 194, 219
  middle-third 194, 219
Cantor space
  disconnected 16
characteristic function 19
  connectedness 19
closed
  uses of the word 13
closed and open sets 10
closed ball
  path-connected 46
closed path 40
closed set
  in a subspace 11
closure
  preserves connectedness 21
compact 109
  discrete metric 110
  finite set 110
  proving compactness 109
  sequentially 76
  set 109
  space 109
compact subsets 109
  finite unions 121
  intersection of 120
compactness 109
complete metric space
  $(C[0,1], d_{\max})$ 151
  $(\mathbf{C}, d_{\mathbf{C}})$ 151
  closed subsets of 147
  $(\mathbb{R}, d^{(1)})$ 146
  $(\mathbb{R}^n, d^{(n)})$ 150
complete subspaces are closed 147
completeness 155
  is a metric invariant 157
  is not a homeomorphism invariant 155
completion 171
component 22
components
  closed 27
  existence of 26
  not open 26
composition of a map with itself 163
connected 33
  as a subspace 20
  but not path-connected 52
  C[0,1] 48
  empty set 15
  equivalent definitions 17
  Euclidean space 30
  interval 33
  path 43
  product space 36
  set 14
  singleton set 15
  space 14
  subset 20
  topological property 18
connectedness
  characteristic function 19
  preserved by closure 21
continuous
  closed set definition 6
  equivalent characterisations 13
continuous functions
  discrete space 7
contraction mapping 161, 215
Contraction Mapping Theorem 158, 160, 163, 207
  Applications 167
contraction ratio 161, 207
cover
  open 103
  open with no finite subcover 106

(d,e)-continuous 6
(d,e)-uniformly continuous 89
definition
  convergent sequence in a metric space 140
  d-convergent 140
  sequence in a metric space 71
dense subspace 171
d_{H}, Hausdorff distance 212
differential equations
  solving 169

dimension 193, 225, 227, 228
disconnected
  Cantor space 16
  discrete metric 16
  equivalent definitions 17
  irrationals 17
  rationals 17
  set 14
  space 14
disconnection 14
discrete metric
  disconnected 16
discrete space
  continuous functions 7
distance
  Euclidean 208, 210
  Hausdorff 212
$d_{\min}$ 114
d-open cover 103

ε-δ definition of continuity 87
empty set
  compact 110
  connected 15
  path-connected 43
ε-net 111
equicontinuous set of functions 91
equivalent metrics
  homeomorphic 7
Euclidean distance
  between two compact sets 210
  from a point to a compact set 208
Euclidean space
  connectedness 30
Extreme Value Theorem 69, 75
  General 84

final point 41
finite subcover 104
fixed point 158
fractal 193

G_d, grid 226
greatest lower bound 30

Hausdorff distance 212
Heine–Borel Theorem
  compact version for $\mathbb{R}^n$ 120
  sequentially compact version for $\mathbb{R}$ 77
  sequentially compact version for $\mathbb{R}^n$ 83
homeomorphic 4
homeomorphism 4
  restriction rule 4

incomplete metric space
  $C[0,1]$ with integral metric 153
infimum 30
initial point 41
integral equation 160
Intermediate Value Theorem 34
interval 8, 33
  classes 10, 35
  path-connected 44
invariant
  topological 8
invariant (for mappings) 218
irrationals
  disconnected 17
isometric 156
isometric invariant 172
isometry 156
iterated function scheme 217

K_{C} 195
K_{G} 200
$\mathcal{K}(\mathbb{R}^m)$ 208
  completeness 213
K_{V} 203, 219

least upper bound 30
  property 30
Lebesgue integrable functions 179
Lebesgue integral 179
Lebesgue number 117
Lipschitz
  constant 161
  function 161

metric space
  compact 109
  complete 146
  completion of 172
  connected 14
  disconnected 14
  incomplete 146
  sequentially compact 76
middle-third Cantor set 194, 219
  totally disconnected 29
min metric 114
monotonic
  decreasing sequence in $\mathbb{R}$ 71
  increasing sequence in $\mathbb{R}$ 71

n-fold composition, T^n 163
$N_d(A)$ 226

open ball
  path-connected 46

open cover 103
    no finite subcover 106
open set
    open ball characterisation 102
open set condition 230
    and self-similar sets 231, 235
open sets
    in a subspace 11

path 40
    closed 41
    final point 41
    initial point 41
path-connected
    C[0,1] 47
    empty set 43
    interval 44
    open and closed balls 46
    $\mathbb{R}$ 44
    $\mathbb{R}^n$ 48
    real space 45
    set 43
    singleton set 43
    space 43
    topological invariant 45
    unit ball 46
path-connectedness
    of products 49
    topological invariant 45
pointwise bounded set of functions 90
products
    path-connectedness 49
    sequentially compact 80
proper subset 15

$\mathbb{R}$
    complete 145
    path-connected 44
$\mathbb{R}^n$
    complete 150
    path-connected 48
rabbit
    hat, out of 9
rationals
    disconnected 17
real sequence
    bounded 119
restriction rule
    homeomorphisms 4

self-similar 198
    and the open set condition 231, 235
sequence 71
    Cauchy 142
    convergent 140
    divergent 140
    limit 140
sequence in a metric space
    definition 71
sequences in $\mathbb{R}$
    monotonic decreasing 71
    monotonic increasing 71
sequentially compact 76
    finite sets 78
    products 80
    subset 76
sequentially compact subsets
    finite unions 122
    intersection of 120
set
    bounded 81
    bounded, ball definition 82
set of functions
    equicontinuous 91
    pointwise bounded 90
Sierpiński gasket 199, 200, 222, 236
    construction of 199, 222
similarity 197
singleton set
    connected 15
    path-connected 43
subcover
    finite 104
subsequence 72
subset
    compact 109
    connected 20
    sequentially compact 76
subspace
    connected 20
    open sets 11
supremum 30

Tikhonov's Theorem 80, 83
topological invariant 8
    path-connectedness 45
topological property
    connectedness 18
topologist's cosine 52
totally bounded 115
totally disconnected 22
    middle-third Cantor set 29
totally disconnected space 29
trail 41

uniformisation argument 153

unit ball
    path-connected 46

von Koch curve 202, 218, 236
    construction of 202, 218

von Koch snowflake 193

zeros of real functions 167